Strategic Delay and Information Exchange in Endogenous Social Networks

by

Kostas Bimpikis

B.Sc., National Technical University of Athens, Greece (2003)
M.Sc., University of California, San Diego (2005)

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2010

Author ....
Sloan School of Management
August 11, 2010

Certified by ....
Daron Acemoglu
Professor
Thesis Supervisor

Certified by ....
Asuman Ozdaglar
Associate Professor
Thesis Supervisor

Accepted by ....
Dimitris Bertsimas
Co-director, Operations Research Center

Report Documentation Page

Report Documentation Page	Form Approved OMB No. 0704-0188

1. REPORT DATE **SEP 2010**	2. REPORT TYPE	3. DATES COVERED **00-00-2010 to 00-00-2010**
4. TITLE AND SUBTITLE **Strategic Delay and Information Exchange in Endogenous Social Networks**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Massachusetts Institute of Technology,Cambridge,MA,02139**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**		
13. SUPPLEMENTARY NOTES		

14. ABSTRACT

This thesis studies optimal stopping problems for strategic agents in the context of two economic applications: experimentation in a competitive market and information exchange in social networks. The economic agents (firms in the first application, individuals in the second) take actions, whose payoffs depend on an unknown underlying state. Our framework is characterized by the following key feature: agents time their actions to take advantage of either the outcome of the actions of others (experimentation model) or information obtained over time by their peers (information exchange model). Equilibria in both environments are typically inefficient, since information is imperfect and, thus, there is a benefit in being a late mover, but delaying is costly. More specifically, in the first part of the thesis, we develop a model of experimentation and innovation in a competitive multi-firm environment. Each firm receives a private signal on the success probability of a research project and decides when and which project to implement. A successful innovation can be copied by other firms. We start the analysis by considering the symmetric environment, where the signal quality is the same for all firms. Symmetric equilibria (where actions do not depend on the identity of the firm) always involve delayed and staggered experimentation, whereas the optimal allocation never involves delays and may involve simultaneous rather than staggered experimentation. The social cost of insufficient experimentation can be arbitrarily large. Then we study the role of simple instruments in improving over equilibrium outcomes. We show that appropriately-designed patents can implement the socially optimal allocation (in all equilibria) by encouraging rapid experimentation and efficient ex post transfer of knowledge across firms. In contrast to patents, subsidies to experimentation, research, or innovation cannot typically achieve this objective. We also discuss the case when signal quality is private information and differs across firms. We show that in this more general environment patents again encourage experimentation and reduce delays. In the second part, we study a model of information exchange among rational individuals through communication and investigate its implications for information aggregation in large societies. An underlying state (of the world) determines which action has higher payoff. Agents receive a private signal correlated with the underlying state. They then exchange information over their social network until taking an (irreversible) action. We define asymptotic learning as the fraction of agents taking an action that is close to optimal converging to one in probability as a society grows large. Under truthful communication

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**165**	

Strategic Delay and Information Exchange in Endogenous Social Networks

by

Kostas Bimpikis

Submitted to the Sloan School of Management
on August 11, 2010, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

This thesis studies optimal stopping problems for strategic agents in the context of two economic applications: experimentation in a competitive market and information exchange in social networks. The economic agents (firms in the first application, individuals in the second) take actions, whose payoffs depend on an unknown underlying state. Our framework is characterized by the following key feature: agents time their actions to take advantage of either the outcome of the actions of others (experimentation model) or information obtained over time by their peers (information exchange model). Equilibria in both environments are typically inefficient, since information is imperfect and, thus, there is a benefit in being a late mover, but delaying is costly.

More specifically, in the first part of the thesis, we develop a model of experimentation and innovation in a competitive multi-firm environment. Each firm receives a private signal on the success probability of a research project and decides when and which project to implement. A successful innovation can be copied by other firms. We start the analysis by considering the symmetric environment, where the signal quality is the same for all firms. Symmetric equilibria (where actions do not depend on the identity of the firm) always involve delayed and staggered experimentation, whereas the optimal allocation never involves delays and may involve simultaneous rather than staggered experimentation. The social cost of insufficient experimentation can be arbitrarily large. Then, we study the role of simple instruments in improving over equilibrium outcomes. We show that appropriately-designed *patents* can implement the socially optimal allocation (in all equilibria) by encouraging rapid experimentation and efficient *ex post* transfer of knowledge across firms. In contrast to patents, subsidies to experimentation, research, or innovation cannot typically achieve this objective. We also discuss the case when signal quality is private information and differs across firms. We show that in this more general environment patents again encourage experimentation and reduce delays.

In the second part, we study a model of information exchange among rational individuals through communication and investigate its implications for information aggregation in large societies. An *underlying state* (of the world) determines which action has higher

payoff. Agents receive a *private signal* correlated with the underlying state. They then exchange information over their *social network* until taking an (irreversible) action. We define *asymptotic learning* as the fraction of agents taking an action that is close to optimal converging to one in probability as a society grows large. Under truthful communication, we show that asymptotic learning occurs if (and under some additional conditions, also only if) in the social network most agents are a short distance away from "information hubs", which receive and distribute a large amount of information. Asymptotic learning therefore requires information to be aggregated in the hands of a few agents. We also show that while truthful communication is not always optimal, when the communication network induces asymptotic learning (in a large society), truthful communication is an equilibrium. Then, we discuss the welfare implications of equilibrium behavior. In particular, we compare the aggregate welfare at equilibrium with that of the *optimal allocation*, which is defined as the strategy profile a social planner would choose, so as to maximize the expected aggregate welfare. We show that when asymptotic learning occurs all equilibria are efficient. A partial converse is also true: if asymptotic learning does not occur at the optimal allocation and an additional mild condition holds at an equilibrium, then the equilibrium is inefficient. Furthermore, we discuss how our learning results can be applied to several commonly studied random graph models, such as preferential attachment and Erdős-Renyi graphs.

In the final part, we study strategic network formation in the context of information exchange. In particular, we relax the assumption that the social network over which agents communicate is fixed, and we let agents decide which agents to form a *communication link* with incurring an associated cost. We provide a systematic investigation of what types of cost structures and associated social cliques (consisting of groups of individuals linked to each other at zero cost, such as friendship networks) ensure the emergence of communication networks that lead to asymptotic learning. Our result shows that societies with too many and sufficiently large social cliques do not induce asymptotic learning, because each social clique would have sufficient information by itself, making communication with others relatively unattractive. Asymptotic learning results if social cliques are neither too numerous nor too large, in which case communication across cliques is encouraged.

Thesis Supervisor: Daron Acemoglu
Title: Professor

Thesis Supervisor: Asuman Ozdaglar
Title: Associate Professor

Acknowledgments

First and foremost, I would like to express my sincerest gratitude to my fantastic advisors, Daron Acemoglu and Asu Ozdaglar. I will be eternally indebted to them for their immense efforts in providing all the essentials for my development as an academic researcher.

Collaborating with Daron is a continuous educational experience, since he is not only an exceptional researcher but also a unique mentor. I learned a great deal from him on how to identify promising research directions, develop interesting economic models and form my own personal research identity. Apart from his astonishing knowledge and mentorship, he is simply a great person to work with: he always treated me as an equal and he was constantly encouraging even in my most unproductive days.

There are no words to describe how fortunate I was to have Asu as my academic advisor. She is one of the most motivated people I have ever met and she is capable of transmitting her enthusiasm to all her collaborators. Her attention to detail without sacrificing the big picture is absolutely stunning and in retrospect I understand the importance of the countless remarks she made on my early research drafts.

In addition to our research collaboration, I was privileged to serve as Asu's and Daron's teaching assistant. Although their insistence on perfection meant more work on my side, I was rewarded by their well-prepared and truly engaging lectures, which helped me mature as a lecturer. Finally, I am much obliged for their great support during my academic job search, which they experienced as if it was their own.

Furthermore, I would like to thank the other two members of my thesis committee, Georgia Perakis and John Tsitsiklis for their careful reading of the thesis and thorough comments. I am also extremely grateful to Georgia for her help and guidance during my academic job search.

Moreover, I would like to extend my gratitude to the ORC and LIDS staff, Paulette

Mosley, Laura Rose, Andrew Carvalho and Lynne Dell, for their help on various administrative tasks. My thesis work was financially supported by the US Air Force and Army Office of Scientific Research, and NSF.

I was extremely fortunate to call both ORC and LIDS my academic homes. My friends and colleagues made my PhD experience smooth and intellectually stimulating. First, I would like to thank my classmates from the ORC, Rajiv Menjoge, Juliane Dunkel, Doug Fearing, Pallav Chhaochharia, Andy Sun, Karima Nigmatulina and Ruben Lobel, as well as the rest of the ORC crowd, Theo Weber, Hamed Mamani, Margret Vilborg, Katy Kaminski, Sun Wei and Nikolas Trichakis. Dan Iancu and Doug Fearing shared with me all the difficulties of the academic job search and they were constantly helpful and encouraging. Second, I am much indebted to the members of Asu's group, Paul Njoroge, Ishai Menache, Ermin Wei, Ali ParandehGheibi, Noah Stein, Alireza Tahbaz-Salehi, Diego Feijer, Kimon Drakopoulos and especially Ozan Candogan and Ilan Lobel, my good friends and collaborators. A special thanks goes to Joline Uichanco, who helped me go through the emotional stress of my last two PhD years (and was the creative mind behind the figures in my job market presentation).

Cambridge would not be as enjoyable if it was not for the company of Giorgos Papachristoudis, Ioannis Simaiakis, Gerasimos Skaltsas, Mohamed Mostagir, Gerry Tsoukalas, Christina Sgardeli, Despina Zymnis, Antonis Vytiniotis, Ioannis Bertsatos and, of course, my dearest friends and roommates, Nikolas Pyrgiotis and Ruben Lobel. Living in the "church" is one of the most enjoyable periods in my life so far.

Finally, I would like to express my deepest appreciation to my parents, Yannis and Magda, who are always there for me. This thesis is dedicated to them.

Contents

List of Figures

Chapter 1

Introduction

Optimal stopping problems have been considered extensively in the literature. They typically involve a single agent or a group of agents that are faced with the problem of deciding when to take an action, so as to maximize their utility. The decision problem is non-trivial due to the following tradeoff: delaying leads to a "better" action, however, future is discounted and, thus, agents cannot wait forever. Moreover, the optimal stopping time for an agent depends in many cases crucially on the action profiles of her peers, thus necessitating incorporating strategic considerations into the framework and treating it as a dynamic game.

In this thesis, we formulate *optimal stopping games* or *waiting games* (see [22]) to study two seemingly unrelated economic applications: (1) investment in R&D in the presence of competitors, that may choose to free-ride on one's experimentation efforts and (2) information exchange among rational individuals, that are embedded in a social network structure. We provide a characterization of the equilibria in both settings and discuss conditions under which they are *inefficient*, i.e., there exist allocations, that achieve a higher expected aggregate welfare.

Alongside with the characterization of equilibria and discussion of welfare loss, we provide a number of results that are specific to each of the two applications. In particular,

in the first model (investment in R&D) we investigate the role of simple policies for improving the aggregate welfare. We show that although subsidies from third parties fail to restore the distorted incentives, simple patents may succeed (under certain conditions) in both dealing with costly delay and allowing firms to share their knowledge. In the second environment, we are interested in whether agents eventually *learn*, i.e., choose an action, that is close to the optimal, and when both objectives (learning and maximizing welfare) coincide.

Roughly the thesis is decomposed into three parts: the first (Chapters 2-4) discusses our model of experimentation and innovation, the second (Chapters 5-7) studies information exchange in already formed social networks and, finally, the last part (Chapter 8) deals with the complementary problem of network formation: given that agents are interested in acquiring as much information as early as possible, what kinds of communication networks arise and what are their learning properties? The following two sections describe our models in a high level and discuss relevant literature. Finally, in Section 1.3 we provide a more detailed roadmap of the thesis.

1.1 Strategic delay in a model of experimentation

Most modern societies provide intellectual property rights protection to innovators using a patent system. The main argument in favor of patents is that they encourage *ex ante* innovation by creating *ex post* monopoly rents (e.g., [4], [7], [68], [34], [35], [39], [53], [54], [58], [59], [60]). In the first part of the thesis, we suggest an alternative (and complementary) social benefit to patents. We show that, under certain circumstances, patents encourage experimentation by potential innovators while still allowing socially beneficial transmission of knowledge across firms.

We construct a stylized model of experimentation and innovation. In our baseline game, N symmetric potential innovators (firms) have each access to a distinct research opportunity and a private signal on how likely it is to result in a successful innovation. Firms can

decide to experiment at any point in time. A successful innovation is publicly observed and can be copied by any of the other potential innovators (for example, other firms can build on the knowledge revealed by the innovation in order to increase their own probability of success, but in the process capture some of the rents of this first innovator). The returns from the implementation of a successful innovation are nonincreasing in the number of firms implementing it. We provide an explicit characterization of the equilibria of this dynamic game. The symmetric equilibrium always features delayed and staggered experimentation. In particular, experimentation does not take place immediately and involves one firm experimenting before others (and the latter firms free-riding on the former's experimentation). In contrast, the optimal allocation never involves delays and may require simultaneous rather than staggered experimentation. The insufficient equilibrium incentives for experimentation may create a significant efficiency loss: the ratio of social surplus generated by the equilibrium relative to the optimal allocation can be arbitrarily small.

We next show that a simple arrangement resembling a patent system, where a copying firm has to make a prespecifed payment to the innovator, can implement the optimal allocation (we refer to this arrangement as a "patent system"). When the optimal allocation involves simultaneous experimentation, the patent system makes free-riding prohibitively costly and implements the optimal allocation as the unique equilibrium. When the optimal allocation involves staggered experimentation, the patent system plays a more subtle role. It permits *ex post* transmission of knowledge but still increases experimentation incentives to avoid delays. The patent system can achieve this because it generates "conditional" transfers. An innovator receives a patent payment only when copied by other firms. Consequently, patents encourage one firm to experiment earlier than others, thus achieving rapid experimentation without sacrificing useful transfer of knowledge. Moreover, we show that patents can achieve this outcome in all equilibria. The fact that patents are particularly well designed to play this role is also highlighted by our result that while an appropriately-designed patent implements the optimal allocation in all equilibria, subsidies

to experimentation, research, or innovation cannot achieve the same objective.

In our baseline model, both the optimal allocation and the symmetric equilibrium involve sequential experimentation. Inefficiency results from lack of sufficient experimentation or from delays. The structure of equilibria is richer when the strength (quality) of the signals received by potential innovators differs and is also private information. In this case, those with sufficiently strong signals will prefer not to copy successful innovations. We show that in this more general environment patents are again potentially useful (though they cannot typically achieve the optimal allocation).

Although our analysis is purely theoretical, we believe that the insights it generates, in particular regarding the role of patents in discouraging delay in research and encouraging experimentation, are relevant for thinking about the role of the patent systems in practice. Two assumptions are important for our results. The first is that pursuing an unsuccessful research line makes a subsequent switch to an alternative line more costly. We impose this feature in the simplest possible way, assuming that such a switch is not possible, though our qualitative results would not be affected if switching were feasible but costly. We view this assumption as a reasonable approximation to reality. Commitment of intellectual and financial resources to a specific research line or vision is necessary for success, and once such commitment has been made, changing course is not easy.[1] Our second key assumption is that copying of successful projects is possible (without prohibitive patents) and reduces the returns to original innovators. This assumption also appears quite plausible. Copying of a successful project here should be interpreted more broadly as using the information

[1]For example, in the computer industry, firms such as Digital Equipment Corporation (DEC) that specialized in mainframes found it difficult to make a successful switch to personal computers (e.g., [26]). Similarly, early innovators in the cell phone industry, such as Nokia and Ericsson, appear to be slow in switching to the new generation of more advanced wireless devices and smartphones, and have been generally lagging behind companies such as Apple and RIM. Another interesting example comes from the satellite launches. The early technology choice for launching satellites into space relied on large ground-based rockets; despite evidence that using smaller rockets and carrying these to the upper atmosphere using existing aerospace equipment would be considerably cheaper and more flexible, organizations such as NASA have not switched to this new technology, while private space technology companies have (see [43]).

revealed by successful innovation or experimentation, so it does not need to correspond to replicating the exact same innovation (or product),[2] and naturally such copying will have some negative effect on the returns of the original innovator.

In addition to the literature on patents mentioned above, a number of other works are related to the model presented here. First, ours is a simple model of (social) experimentation and shares a number of common features with recent work in this area (e.g., [17] and [52]). These papers characterize equilibria of multi-agent two-armed bandit problems and show that there may be insufficient experimentation. The structure of the equilibrium is particularly simple in our model and can be characterized explicitly because all payoff-relevant uncertainty is revealed after a single successful experimentation. In addition, as discussed above, there is insufficient experimentation in our model as well, though this also takes a simple form: either there is free-riding by some firms reducing the amount of experimentation or experimentation is delayed. We also show that patent systems can increase experimentation incentives and implement the optimal allocation.

Second, the structure of equilibria with symmetric firms is reminiscent to equilibria in war of attrition games (e.g., [42], [46], [64]). War of attrition games have been used in various application domains, such as the study of market exit ([18] and [29]), research tournaments ([67]), auctions ([19]), investment choices ([20]), exploratory drilling ([44] and [45]) and the diffusion of new technologies ([51]). Similar in spirit with our work, [22] discusses *waiting games* of technological change, in which there is a late-mover advantage due to knowledge spillovers. In our symmetric model, as in symmetric wars of attrition, players choose the stochastic timing of their actions in such a way as to make other players indifferent and willing to mix over the timing of their own actions. The structure of equilibria and the optimal allocation is different, however, and the optimal allocation may involve either simultaneous experimentation by all players or staggered experimentation

[2]In terms of the examples in footnote 1, while DEC, Nokia and Ericsson may have been slow in adopting new technologies, several other, new companies have built on the technological advances that took place in personal computers and smartphones.

similar to that resulting in asymmetric equilibria. The novel beneficial role of patents in our model arises from their ability to implement such asymmetric equilibria.

Finally, the monotonicity property when the quality of signals differs across agents is similar to results in generalized wars of attrition (e.g., [18], [19] and [29]) and is also related to [40] result on the clustering of actions in herding models. In the context of a standard herding model with endogenous timing, Gul and Lundholm construct an equilibrium in which agents with stronger signals act earlier than those with weaker signals, though the specifics of our model and analysis differs from these previous contributions. [3]

1.2 Information exchange in endogenous social networks

Most social decisions, ranging from product and occupational choices to voting and political behavior, rely on information agents gather through communication with friends, neighbors and co-workers as well as information obtained from news sources and prominent webpages. A central question in social sciences concerns the dynamics of communication and information exchange and whether such dynamics lead to the effective aggregation of dispersed information that exists in a society. In the second part of the thesis, we construct a dynamic model to investigate this question. If information exchanges were non-strategic, timeless and costless, all information could be aggregated immediately by simultaneous communication across all agents. Thus, the key ingredient of our approach is dynamic and costly communication.

Our benchmark model features an *underlying state* of the world that determines which action has higher payoff (which is the same for all agents). Because of discounting, earlier actions are preferred to later ones. Each agent receives a *private signal* correlated with this underlying state. In addition, agents communicate with those others with whom

[3] This monotonicity property does not hold in our model when there are more than two firms and the support of signals includes sufficiently strong signals so that some firms prefer not to copy successful experimentations as is shown in Appendix B.

they are connected in their *social network* until they take an irreversible action. Crucially, information acquisition takes time because the "neighbors" of an agent with whom she communicates acquire more information from their own neighbors over time. Information exchange will thus be endogenously limited by two features: the communication network , which allows communication only between connected pairs, and discounting, which encourages agents to take actions before they accumulate sufficient information.

We characterize the equilibria of this communication game and then investigate the structure of these equilibria as the society becomes large (i.e., for a sequence of games). Our main focus is on how well information is aggregated, which we capture with the notion of ϵ, δ-asymptotic learning. We say that there is *ϵ, δ-asymptotic learning* if more than $(1 - \epsilon)$-fraction of the agents take an action that is ϵ-close to optimal with probability at least $1 - \delta$, as the society becomes large. Furthermore, we say that *perfect asymptotic learning* occurs, when the fraction of agents that take a nearly optimal action converges to one in probability as the society grows.

Our analysis proceeds in several stages. *First*, we assume that agents are non-strategic in their communication. Under these assumptions, we provide a necessary and a sufficient condition for ϵ, δ-asymptotic learning under a given equilibrium strategy profile. Intuitively, these conditions require that most of the agents have close access to a sufficient amount of information. We also show a sufficient and (under a mild condition) necessary condition for perfect asymptotic learning, that holds for all equilibria and it only involves the network topology. The learning results are further illustrated by a series of corollaries that identify two types of information hubs as conduits of asymptotic learning. The first are *information mavens*, which have a large in-degree, enabling them to aggregate information. If most agents are close to an information maven, asymptotic learning is guaranteed. The second type of hubs are *social connectors*, which have large out-degree, enabling them to communicate their information to a large number of agents.[4] Social connectors are only

[4]both of these terms are inspired by [36].

useful for asymptotic learning if they are close to mavens, so that they can distribute their information. Thus, asymptotic learning is also obtained if most agents are close to a social connector, who is in turn a short distance away from a maven.

Second, we generalize these results to an environment, in which individuals may misreport their information if they have an incentive to do so. In particular, we show that individuals may in general choose to misreport their information in order to delay the action of their neighbors, thus obtaining more information from them in the future. Nevertheless, we establish that whenever truthful communication leads to perfect asymptotic learning, it is an equilibrium of the strategic communication game to report truthfully. Interestingly, the converse is not necessarily true: strategic communication may lead to asymptotic learning in some special cases in which truthful communication precludes learning.

Our characterization results on asymptotic learning can be seen both as "positive" and "negative". On the one hand, communication structures that do not feature such hubs appear more realistic in the context of social networks and communication between friends, neighbors and co-workers. Indeed, the popular (though not always empirically plausible) random graph models such as preferential attachment and Poisson (Erdős-Renyi) graphs do not lead to asymptotic learning. On the other hand, as discussed above, most individuals obtain key information from either individuals or news sources (websites) that correspond to mavens and social connectors, which do play the role of aggregating and distributing large amounts of information. Corresponding to such structures, we show that scale free random graphs (in particular, power law graphs with small exponent $\gamma \leq 2$),[5] and hierarchical graphs, where "special" agents are likely to receive and distribute information to lower layers of the hierarchy, induce network structures that guarantee asymptotic learning. The intuition for why such information hubs and almost all agents being close to information hubs are necessary for asymptotic learning is instructive: were it not so, a large fraction

[5] These models are shown to provide good representations for peer-to-peer networks, scientific collaboration networks (in experimental physics), and traffic in networks ([50], [57], [69] and [70].)

of agents would prefer to take an action before waiting for sufficient information to arrive and a nontrivial fraction of those would take a suboptimal action.

Furthermore, we discuss the welfare implications of equilibrium behavior. In particular, we compare the expected aggregate welfare at equilibrium with that at the *optimal allocation*, which is defined as the timing profile that a social planner would choose, so as to maximize it. We show that when perfect asymptotic learning occurs, all equilibria are efficient. Moreover, a partial converse holds: if ϵ, δ-asymptotic learning fails at equilibrium and an additional condition holds, then the equilibrium is inefficient. Our results regarding welfare uncover an interesting feature of communication in social networks: information exchange is inefficient, because agents do not internalize the positive externality that their presence exerts on their peers.

In the third part of the thesis, we drop the assumption that the social network is given and we let agents form *communication links*, which may be costly, thus combining social learning with strategic network formation. The network formation decisions of agents induce a *communication graph* for the society, over which communication takes place in the way described above. Armed with the analysis of information exchange over a given communication network, we analyze the equilibria of the two-stage network formation and information exchange game, under the assumption that although forming communication links is costly, there also exist *social cliques*, groups of individuals that are linked to each other at zero cost. These can be thought of as "friendship networks," which are linked for reasons unrelated to information exchange and thus act as conduits of such exchange at low cost. Agents have to pay a cost at the beginning in order to communicate (receive information) from those who are not in their social clique.

Even though network formation games have several equilibria, the structure of our network formation and information exchange game enables us to obtain relatively sharp results on what types of societies will lead to endogenous communication networks that ensure asymptotic learning. In particular, we show that societies with too many (disjoint)

and sufficiently large social cliques induce behavior inconsistent with asymptotic learning. This is because each social clique, which is sufficiently large, would have enough information to make communication with others (from other social cliques) unattractive; the society gets segregated into a very large number of disjoint social cliques not sharing information. In contrast, asymptotic learning obtains in equilibrium if social cliques are neither too numerous nor too large so that it becomes advantageous at least for some members of these cliques to communicate with members of other cliques, forming a structure in which information is shared across (almost) all members of the society.

Our work is related to several strands of literature on social and economic networks. First, it is related to the large and growing literature on *social learning.* Much of this literature focuses on Bayesian models of observational learning, where each individual learns from the actions of others taken in the past. A key impediment to information aggregation in these models is the fact that actions do not reflect all of the information that an individual has and this can induce a pattern reminiscent to a "herd," where individuals ignore their own information and copy the behavior of others (see, for example, [12], [15] and [65], as well as [10], for early contributions, and [3], [13] and [66] for models of Bayesian learning with richer observational structures). While observational learning is important in many situations, a large part of information exchange in practice is through communication.

Several papers in the literature study communication, though typically using non-Bayesian or "myopic" rules (for example, [23] and [37]). A major difficulty faced by these approaches, often precluding Bayesian and dynamic game theoretic analysis of learning in communication networks, is the complexity of updating when individuals share their ex-post beliefs. We overcome this difficulty by adopting a different approach, whereby individuals can directly communicate their signals and there is no restriction on the total "bits" of communication. This leads to a very tractable structure for updating of beliefs and enables us to study perfect Bayesian equilibria of a game of network formation, communication and decision-making. It also reverses one of the main insights of these papers, also shared

by the pioneering contribution to the social learning literature by [10], that the presence of "highly connected" or "influential" agents, or what [10] call a "royal family," acts as a significant impediment to the efficient aggregation of information. On the contrary, in our model the existence of such highly connected agents (information hubs, mavens or connectors) is crucial for the efficient aggregation of information. Moreover, their existence also reduces incentives for non-truthful communication, and is the key input into our result that truthful communication can be an ϵ-equilibrium.

Our analysis of asymptotic learning in large networks also builds on *random graph models*. In particular, we use several tools and results from this literature to characterize the asymptotics of beliefs and information. We also study information aggregation in the popular *preferential attachment* and *Erdős-Renyi* graphs (e.g., [5], [14], [25] and [55]).

Our work is also related to the growing literature on *network formation*, since communication takes place over endogenously formed networks (Chapter 8). Although the network formation literature is large and growing (see, e.g., [11], [47] and [49]), we are not aware of other papers that endogenize the benefits of forming links through the subsequent information exchange. It is also noteworthy that, while network formation games have a large number of equilibria, the simple structure of our model enables us to derive relatively sharp results about environments in which the equilibrium networks will lead to asymptotic learning.

Finally, our model is related to the literature on *strategic communication*, pioneered by the cheap talk framework of [21]. While cheap talk models have been used for the study of information aggregation with one receiver and multiple senders (e.g. [56]) and multiple receivers and single sender (e.g. [27]), most relevant to our model are two recent papers that consider strategic communication over general networks, [32] and [41]. A major difference between these works and ours is that we consider a model where communication is allowed for more than one time period, thus enabling agents to receive information outside their immediate neighborhood (at the cost of a delayed decision) and we also endogenize

the network over which communication takes place. On the other hand, our framework assumes that an agent's action does not directly influence others' payoffs, while such payoff interactions are the central focus of [32] and [41].

1.3 Organization

The thesis is organized as follows. In Chapter 2, we develop a model of experimentation and innovation in a competitive multi-firm environment. Moreover, we provide an explicit characterization of both the asymmetric and symmetric equilibria, when firms are symmetric, i.e., they receive information of equal precision. Chapter 3 characterizes the optimal allocation and shows that the efficiency gap between the symmetric equilibrium and the optimal allocation can be arbitrarily large. The analysis in this chapter also demonstrates that an appropriately-designed patent system can implement the optimal allocation (in all equilibria). Chapter 4 extends the model to an environment, where firms have different signal qualities. In Chapter 5, we introduce a model of information exchange among rational agents, that are embedded in a social network. Chapter 6 discusses conditions under which agents *learn* at equilibrium and also explores the welfare properties of equilibria. In Chapter 7, we apply our learning results to several models of random graphs. Chapter 8 extends the information exchange model by incorporating endogenous network formation. Finally, Chapter 9 concludes the thesis and discusses a number of promising avenues for future research. Appendices A-F contain additional results and some of the proofs omitted from the main text.

Chapter 2

A Model of Strategic Experimentation

The economy in the baseline model consists of two research firms, each maximizing the present discounted value of profits. Time is continuous,[1] and both firms discount the future at a common rate $r > 0$.

Each firm can implement ("experiment with") a distinct project. The success probability of experimentation is $p > 0$. The success or failure of experimentation by a firm is publicly observed. When experimentation is successful, we refer to this as an "innovation".

At time t, a firm can choose one of three possible actions: (1) experiment with a project (in particular, with the project on which the firm has received a positive signal); (2) copy a successful project; (3) wait. Experimentation and copying are irreversible, so that a firm cannot then switch to implement a different project. In the context of research, this captures the fact that commitment of intellectual and financial resources to a specific research line or project is necessary for success. Copying of a successful project can be interpreted more broadly as using the information revealed by successful innovation or experimentation, so

[1] In the Appendix, we present a discrete-time version of the model and formally show that the continuous-time version studied in the main text gives the same economic and mathematical answers as the limit of the discrete-time model, when the time interval $\Delta \to 0$.

it does not need to correspond to the second firm replicating the exact same innovation (or product).[2]

Payoffs depend on the success of the project and whether the project is copied. During an interval of length τ, the payoff to a firm that is the only one implementing a successful project is $\pi_1\tau > 0$. In contrast, if a successful project is implemented by both firms, each receives $\pi_2\tau > 0$.[3] The payoff to an unsuccessful project is normalized to zero.

Until we introduce heterogeneity in success probabilities, we maintain the following assumption.[4]

Assumption 1.

$$\pi_1 > \pi_2 > p\pi_1.$$

Let us also define the present discounted value of profits as

$$\Pi_j \equiv \frac{\pi_j}{r} \text{ for } j = 1, 2,$$

and for future reference, define

$$\beta \equiv \frac{\Pi_2}{\Pi_1}. \tag{2.1}$$

Clearly, $\beta \in (p, 1)$ in view of Assumption 1. Assumption 1 implies that the payoff from a new innovation is decreasing in the number of firms that adopt it ($\pi_1 > \pi_2$) and also that the expected payoff of a firm's experimentation is smaller than the payoff from copying a successful innovation. In particular, the firm prefers to copy than to experiment with its own research opportunity.

[2]Several generalizations do not affect our qualitative results, and are not introduced to reduce notation and maximize transparency. These include: (1) copying can be successful with some probability $\nu \in (p, 1]$; (2) copying after an unsuccessful experimentation is feasible, but involves a cost $\Gamma_1 > 0$; (3) experimentation itself involves a cost $\Gamma_2 > 0$.

[3]It will be evident from the analysis below that all of our results can be straightforwardly generalized to the case where an innovator receives payoff $\pi_2^{\text{first}}\tau$ when copied, whereas the copier receives $\pi_2^{\text{second}}\tau$. Since this has no major effect on the main economic insights and just adds notation, we do not pursue this generalization.

[4]The structure of equilibria without this assumption is trivial as our analysis in Section 3 shows.

Now we are in a position to define strategies in this game. Let a history up to time t be denoted by h^t. The set of histories is denoted by $\mathcal{H}^t$. A strategy for a firm is a mapping from time t and the history of the game up to t, h^t, to the *flow rate of experimentation* at time t and the distribution over projects. Thus, the time t strategy can be written as

$$\sigma : \mathbb{R}_+ \times \mathcal{H}^t \to \bar{\mathbb{R}}_+ \times \boldsymbol{\Delta}(\{1,2\}),$$

where $\bar{\mathbb{R}}_+ \equiv \mathbb{R}_+ \cup \{+\infty\}$ and $\boldsymbol{\Delta}(\{1,2\})$ denotes the set of probability distributions over the set of projects (project available to the first, second firm is labeled 1, 2 respectively), corresponding to the choice of project when the firm implements one. The latter piece of generality is largely unnecessary (and will be omitted), since there will never be mixing over projects (a firm will either copy a successful project or experiment with the project for which it has received a positive signal). Here $\sigma(t, h^t) = (0, \cdot)$ corresponds to waiting at time t and $\sigma(t, h^t) = (\infty, j)$ corresponds to implementing project j at time t, which could be experimentation or copying of a successful project. Let us also denote the strategy of firm $i = 1, 2$ by $\sigma_i = \{\sigma_i(t, .)\}_{t=0}^{\infty}$.

History up to time t can be summarized by two events $a^t \in \{0, 1\}$ denoting whether the other firm has experimented up to time t and $s^t \in \{0, 1\}$ denoting whether this choice was successful. With a slight abuse of notation we will use both $\sigma(t, h^t)$ and $\sigma(t, , a^t, s^t)$ to denote time t strategies. We study subgame perfect equilibria in the environment defined above. In particular, a subgame perfect equilibrium (or simply *equilibrium*) is a strategy profile $(\hat{\sigma}_1, \hat{\sigma}_2)$ such that $(\hat{\sigma}_1|h^k, \hat{\sigma}_2|h^k)$ is a Nash equilibrium of the subgame defined by history h^t for all histories $h^t \in \mathcal{H}^t$, where $\hat{\sigma}_i|h^k$ denotes the restriction of $\hat{\sigma}_i$ to the histories consistent with h^k.

2.1 Asymmetric Equilibria

Even though firms are symmetric (in terms of their payoffs and information), there can be symmetric and asymmetric equilibria. Our main interest is with symmetric equilibria,

where strategies are independent of the identity of the player. Nevertheless, it is convenient to start with asymmetric equilibria. These equilibria are somewhat less natural, because, as we will see, they involve one of the players never experimenting until the other one does. Before describing the equilibria, we introduce some additional notation. In particular, the flow rate of experimentation σ induces a stochastic distribution of "stopping time," which we denote by τ. The stopping time τ designates the probability distribution that experimentation will happen at any time $t \in \bar{\mathbb{R}}_+$ conditional on the other player not having experimented until then. A pure strategy simply specifies $\tau \in \bar{\mathbb{R}}_+$. For example, the strategy of experimenting immediately is $\tau = 0$, whereas that of waiting for the other firm's experimentation is represented by $\tau = +\infty$. The τ notation is convenient to use for the next two propositions, while in characterizing the structure of equilibria we need to use σ (thus justifying the introduction of both notations).[5]

In an asymmetric equilibrium, one of the firms, say 1, experiments immediately with its research project. Firm 2 copies firm 1 immediately afterwards if the latter is successful and tries its own project otherwise. Throughout the paper, when there are two firms, we use the notation $\sim i$ to denote the firm $i' \neq i$.

Proposition 1. *Suppose that Assumption 1 holds. Then there exist two asymmetric equilibria. In each equilibrium,* $\tau_i = 0$ *and* $\tau_{\sim i} = +\infty$ *for* $i = 1, 2$.

The proof of the proposition is straightforward and can be therefore omitted. Note that the equilibria described above are not the only asymmetric equilibria in this environment. Another set of such equilibria involves one of the firms experimenting with positive probability (not going to zero) at time $t = 0$ and both firms using a constant flow of experimentation from then on. The crucial feature of asymmetric equilibria is that they explicitly

[5] Here we could follow the more precise approach in [63] for modeling strategies in continuous-time games with jumps. This amounts to defining an extended strategy space, where stopping times are defined for all $t \in \bar{\mathbb{R}}_+$ and also for $t+$ for any $t \in \mathbb{R}_+$. In other words, strategies will be piecewise continuous and right continuous functions of time, so that a jump immediately following some time $t \in \mathbb{R}_+$ is well defined. Throughout, we do allow such jumps, but do not introduce the additional notation, since this is not necessary for any of the main economic insights or proofs.

condition on the identity of the firms.

2.2 Symmetric Equilibria

As mentioned above, asymmetric equilibria explicitly condition on the identity of the firm: one of the firms, with label i, is treated differently than the firm with label $\sim i$. This has important payoff consequences. In particular, it can be verified easily that firm $\sim i$ has strictly greater payoffs in the equilibrium of Proposition 1 than firm i. In addition, as already noted in the previous section, asymmetric equilibria rely on some degree of coordination between the firms, e.g., one of them will not experiment until the other one does. In this light, symmetric equilibria, where strategies are not conditioned on firms' "labels," and firms obtain the same equilibrium payoffs in expectation are more natural. In this section, we study such symmetric equilibria.

As defined above a firm's strategy is a mapping from time and the firm's information set to the flow rate of experimentation with a project. We refer to a strategy as pure if the flow rate of experimentation at a given time t is either 0 or ∞. Our first result shows that there are no pure-strategy symmetric equilibria.

Proposition 2. *Suppose that Assumption 1 holds. Then there exist no symmetric pure-strategy equilibria.*

Proof. Suppose, to obtain a contradiction, that a symmetric pure-strategy equilibrium exists. Then $\tau^* = t \in \mathbb{R}_+$ for $i = 1, 2$, yielding payoff

$$V(\tau^*, \tau^*) = e^{-rt} p \Pi_1$$

to both players. Now consider a deviation $\tau' > t$ for one of the firms, which involves waiting for a time interval ϵ and copying a successful innovation if there is such an innovation during

this time interval. As $\epsilon \to 0$, this strategy gives the deviating firm payoff equal to

$$V(\tau', \tau^*) = \lim_{\epsilon \downarrow 0} e^{-r(t+\epsilon)} \left[p\Pi_2 + (1-p)p\Pi_1\right].$$

Assumption 1 implies that $V(\tau', \tau^*) > V(\tau^*, \tau^*)$, establishing the result. ■

Proposition 2 is intuitive. Asymmetric equilibria involve one of the firms always waiting for the other one to experiment and receiving higher payoff. Intuitively, Proposition 2 implies that in symmetric equilibria both firms would like to be in the position of the firm receiving higher payoffs and thus delaying their own experimentation in order to benefit from that of the other firm. These incentives imply that no (symmetric) equilibrium can have immediate experimentation with probability 1 by either firm.

Proposition 2 also implies that all symmetric equilibria must involve mixed strategies. Moreover, any candidate equilibrium strategy must involve copying of a successful project in view of Assumption 1 and immediate experimentation when the other firm has experimented. Therefore, we can restrict attention to time t strategies of the form

$$\hat{\sigma}_i\left(t, a^t, s^t\right) = \begin{cases} (1, \sim i) & \text{if } a^t = 1 \text{ and } s^t = 1, \\ (1, i) & \text{if } a^t = 1 \text{ and } s^t = 0, \\ (\lambda(t), i) & \text{if } a^t = 0, \end{cases} \tag{2.2}$$

where $\lambda : \mathbb{R}_+ \to \bar{\mathbb{R}}_+$ designates the flow rate of experimentation at time t conditional on no experimentation by either firm up to time t. Given this observation, from now on we will work directly with $\lambda(t)$.

Next we derive an explicit characterization of the (unique) symmetric equilibrium. The next lemma (proof in the appendix) shows that symmetric equilibria must involve mixing on all $t \in \mathbb{R}_+$ and will be used in the characterization of mixed-strategy equilibria.

Lemma 1. *The support of mixed strategy equilibria is* $\mathbb{R}_+$.

Lemma 1 implies that in all symmetric equilibria there will be mixing at all times (until

there is experimentation). Using this observation, Proposition 3 characterizes the unique symmetric equilibrium. Let us illustrate the reasoning here by assuming that firms use a constant flow rate of experimentation (the proof of Proposition 3 relaxes this assumption). In particular, suppose that firm $\sim i$ innovates at the flow rate λ for all $t \in \mathbb{R}_+$. Then the value of innovating at time t (i.e., choosing $\tau = t$) for firm i is

$$V(t) = \int_0^t \lambda e^{-\lambda z} e^{-rz} \left[p\Pi_2 + (1-p)\, p\Pi_1\right] dz + e^{-\lambda t} e^{-rt} p\Pi_2. \tag{2.3}$$

This expression uses the fact that when firm $\sim i$ is experimenting at the flow rate λ, the timing of its experimentation has an exponential distribution, with density $\lambda e^{-\lambda t}$. Then the first term in (2.3) is the expected discounted value from the experimentation of firm $\sim i$ between 0 and t (again taking into account that following an experimentation, a successful innovation will be copied, with continuation value $p\Pi_2 + (1-p)\, p\Pi_1$). The second term is the probability that firm $\sim i$ does not experiment until t, which, given the exponential distribution, is equal to $e^{-\lambda t}$, multiplied by the expected discounted value to firm i when it is the first to experiment at time t (given by $e^{-rt} p\Pi_2$).

Lemma 1 implies that $V(t)$ must be constant in t for all $t \in \mathbb{R}_+$. Therefore, its derivative $V'(t)$ must be equal to zero for all t in any symmetric equilibrium implying that

$$V'(t) = \lambda e^{-\lambda t} e^{-rt} \left[p\Pi_2 + (1-p)\, p\Pi_1\right] - (r+\lambda)\, e^{-\lambda t} e^{-rt} p\Pi_2 = 0, \tag{2.4}$$

for all t. This equation has a unique solution:

$$\lambda^* \equiv \frac{r\beta}{1-p} \text{ for all } t. \tag{2.5}$$

The next proposition shows that this result also holds when both firms can use time-varying experimentation rates.

Proposition 3. *Suppose that Assumption 1 holds. Then there exists a unique symmetric*

equilibrium. This equilibrium involves both firms using a constant flow rate of experimentation λ^* as given by (2.5). Firm i immediately copies a successful innovation by firm $\sim i$ and experiments if firm $\sim i$ experiments unsuccessfully.

Proof. Suppose that firm $\sim i$ experiments at the flow rate $\lambda(t)$ at time $t \in \mathbb{R}_+$. Let us define

$$m(t) \equiv \int_0^t \lambda(z)dz. \tag{2.6}$$

Then the equivalent of (2.3) is

$$V(t) = \int_0^t \lambda(z) e^{-m(z)} e^{-rz} [p\Pi_2 + (1-p) p\Pi_1] dz + e^{-m(t)} e^{-rt} p\Pi_2. \tag{2.7}$$

Here $\int_{t_1}^{t_2} \lambda(z)e^{-m(z)}dz$ is the probability that firm $\sim i$ (using strategy λ) will experiment between times t_1 and t_2, and $e^{-m(t)} = 1 - \int_0^t \lambda(z) e^{-m(z)}dz$ is the probability that $\sim i$ has not experimented before time t. Thus the first term is the expected discounted value from the experimentation of firm $\sim i$ between 0 and t (discounted and multiplied by the probability of this event). The second term is again the probability that firm $\sim i$ does not experiment until t multiplied by the expected discounted value to firm i when it is the first to experiment at time t (given by $e^{-rt}p\Pi_2$).

Lemma 1 implies that $V(t)$ must be constant in t for all $t \in \mathbb{R}_+$. Since $V(t)$ is differentiable in t, this implies that its derivative $V'(t)$ must be equal to zero for all t. Therefore,

$$\begin{aligned} V'(t) &= \lambda(t) e^{-m(t)} e^{-rt} [p\Pi_2 + (1-p) p\Pi_1] - (r + m'(t)) e^{-m(t)} e^{-rt} p\Pi_2 \\ &= 0 \text{ for all } t. \end{aligned}$$

Moreover, note that $m(t)$ is differentiable and $m'(t) = \lambda(t)$. Therefore, this equation is equivalent to

$$\lambda(t) [p\Pi_2 + (1-p) p\Pi_1] = (r + \lambda(t)) p\Pi_2 \text{ for all } t. \tag{2.8}$$

The unique solution to (2.8) is (2.5), establishing the uniqueness of the symmetric equilibrium without restricting strategies to constant flow rates. ■

We end this section by discussing how the equilibrium flow rate of experimentation λ^*, given by equation (2.5), is affected by the relevant parameters. In particular, consider increasing π_1 (the inequality $\pi_2 > p \cdot \pi_1$ would clearly continue to hold). This increases the value of waiting for a firm and leaves the value of experimenting unchanged, so the equilibrium flow rate of experimentation declines, i.e., increasing π_1 reduces β and λ^*.

2.3 Multiple Firms

Let us now suppose that there are N firms, each of which receives a positive signal about one of the projects. The probability that the project that has received a positive signal will succeed is still p and each firm receives a signal about a different project. Let π_n denote the flow payoff from a project that is implemented by n other firms and define

$$\Pi_n \equiv \frac{\pi_n}{r}.$$

Once again, $\beta \equiv \Pi_2/\Pi_1$ as specified in (2.1) and Assumption 1 holds, so that $\beta > p$.

The following proposition is established using similar arguments to those in the previous two sections and its proof is omitted.

Proposition 4. *Suppose that Assumption 1 holds and that there are $N \geq 2$ firms. Then there exist no symmetric pure-strategy equilibria. Moreover the support of the mixed-strategy equilibria is $\mathbb{R}_+$.*

It is also straightforward to show that there exist asymmetric pure-strategy equilibria. For example, when $\Pi_N/\Pi_1 > p$, it is an equilibrium for firm 1 to experiment and the remaining $N-1$ to copy if this firm is successful. If it is unsuccessful, then firm 2 experiments and so on.

As in the previous two sections, symmetric equilibria are of greater interest. To char-

acterize the structure of symmetric equilibria, let us first suppose that

$$\Pi_n = \Pi_2 \text{ for all } n \geq 2 \tag{2.9}$$

and also to simplify the discussion, focus on symmetric equilibria with constant flow rates. In particular, let the rate of experimentation when there are $n \geq 2$ firms be λ_n. Consider a subgame starting at time t_0 with n firms that have not yet experimented (and all previous, $N - n$, experiments have been unsuccessful). Then the continuation value of firm i (from time t_0 onwards) when it chooses to experiment with probability 1 at time $t_0 + t$ is

$$v_n(t) = \int_{t_0}^{t_0+t} \lambda_n (n-1) e^{-\lambda_n(n-1)(z-t_0)} e^{-r(z-t_0)} \left[p\Pi_2 + (1-p) v_{n-1}\right] dz + e^{-\lambda_n(n-1)t} e^{-rt} p\Pi_2, \tag{2.10}$$

where v_{n-1} is the maximum value that the firm can obtain when there are $n-1$ firms that have not yet experimented (where we again use v since this expression refers to the continuation value from time t_0 onwards). Intuitively, $\lambda_n (n-1) e^{-\lambda_n(n-1)(z-t_0)}$ is the density at which one of the $n-1$ other firms mixing at the rate λ_n will experiment at time $z \in (t_0, t_0 + t)$. When this happens, it is successful with probability p and will be copied by all other firms, and each will receive a value of $e^{-r(z-t_0)}\Pi_2$ (discounted back to t_0). If it is is unsuccessful (probability $1-p$), the number of remaining firms is $n-1$, and this gives a value of v_{n-1}. If no firm experiments until time t, firm i chooses to experiment at this point and receives $e^{-rt}p\Pi_2$. The probability of this event is $1 - \int_{t_0}^{t_0+t} \lambda_n (n-1) e^{-\lambda_n(n-1)(z-t_0)} dz = e^{-\lambda_n(n-1)t}$. As usual, in a mixed strategy equilibrium, $v_n(t)$ needs to be independent of t and moreover, it is clearly differentiable in t. So its derivative must be equal to zero. This implies

$$\lambda_n (n-1) \left[p\Pi_2 + (1-p) v_{n-1}\right] = \left(\lambda_n (n-1) + r\right) p\Pi_2. \tag{2.11}$$

Proposition 4 implies that there has to be mixing in all histories, thus

$$v_n = p\Pi_2 \text{ for all } n \geq 2. \tag{2.12}$$

Intuitively, mixing implies that the firm is indifferent between experimentation and waiting, and thus its continuation payoff must be the same as the payoff from experimenting immediately, which is $p\Pi_2$. Combining (2.11) and (2.12) yields

$$\lambda_n = \frac{r\Pi_2}{(1-p)(n-1)\Pi_1}. \tag{2.13}$$

This derivation implies that in the economy with N firms, each firm starts mixing at the flow rate λ_N. Following an unsuccessful experimentation, they increase their flow rate of experimentation to λ_{N-1}, and so on.

The derivation leading up to (2.13) easily generalizes when we relax (2.9). To demonstrate this, let us relax (2.9) and instead strengthen Assumption 1 to:

Assumption 2.

$$\Pi_n > p\Pi_1 \textit{ for all } n.$$

The value of experimenting at time t (starting with n firms) is now given by a generalization of (2.10):

$$v_n(t) = \int_{t_0}^{t_0+t} \lambda_n (n-1) e^{-\lambda_n(n-1)(z-t_0)} e^{-r(z-t_0)} \left[p\Pi_n + (1-p) v_{n-1}\right] dz + e^{-\lambda_n(n-1)t} e^{-rt} p\Pi_n.$$

Again differentiating this expression with respect to t and setting the derivative equal to zero gives the equivalent indifference condition to (2.11) as

$$\lambda_n (n-1) \left[p\Pi_n + (1-p) v_{n-1}\right] = (\lambda_n (n-1) + r) p\Pi_n. \tag{2.14}$$

for $n = 2, ..., N$. In addition, we still have

$$v_n = p\Pi_n \text{ for all } n \geq 2.$$

Combining this with (2.14), we obtain

$$(n-1)\,\lambda_n = \frac{r}{1-p} \cdot \frac{\Pi_n}{\Pi_{n-1}} \text{ for } n = 2, ..., N, \tag{2.15}$$

and let us adopt the convention that $\lambda_1 = +\infty$. Note that the expression on the left hand side is the aggregate rate of experimentation that a firm is facing from the remaining firms. This derivation establishes the following proposition.

Proposition 5. *Suppose that Assumption 2 holds. Then there exists a unique symmetric equilibrium. In this equilibrium, when there are $n = 1, 2, ..., N$ firms that have not yet experimented, each experiments at the constant flow rate λ_n as given by (2.15). A successful innovation is immediately copied by all remaining firms. An unsuccessful experimentation starting with $n \geq 3$ firms is followed by all remaining firms experimenting at the flow rate λ_{n-1}.*

An interesting feature of Proposition 5 is that after an unsuccessful experimentation, the probability of further experimentation may decline. Whether this is the case or not depends on how fast Π_n decreases in n.

Chapter 3

Patents and Optimal Allocations

The analysis so far has established that symmetric equilibria involve mixed strategies, potential delays, and also staggered experimentation (meaning that with probability 1, one of the firms will experiment before others). Asymmetric equilibria avoid delays, but also feature staggered experimentation. Moreover, they are less natural, because they involve one of the firms never acting (experimenting) until the other one does and also because they give potentially very different payoffs to different firms. In this section, we first establish the inefficiency of (symmetric) equilibria. We then suggest that an appropriately-designed patent system can implement optimal allocations. While all of the results in the section hold for $N \geq 2$ firms, we focus on the case with two firms to simplify notation.

3.1 Welfare

It is straightforward to see that symmetric equilibria are Pareto suboptimal. Suppose that there exists a social planner that can decide the experimentation time for each firm. Suppose also that the social planner would like to maximize the sum of the present discounted values of the two firms. Clearly, in practice an optimal allocation (and thus the objective function of the social planner) may also take into account the implications of these innovations on consumers. However, we have not so far specified how consumer welfare is affected by the replication of successful innovations versus new innovations. Therefore, in

what follows, we focus on optimal allocations from the viewpoint of firms. This would also be the optimal allocation taking consumer welfare into account when consumer surpluses from a new innovation and from a successful innovation implemented by two firms are proportional to Π_1 and $2\Pi_2$, respectively. If we take the differential consumer surpluses created by these innovations into account, this would only affect the thresholds provided below, and for completeness, we also indicate what these alternative thresholds would be.

The social planner could adopt one of two strategies:

1. *Staggered experimentation:* this would involve having one of the firms experiment at $t = 0$; if it is successful, then the other firm would copy the innovation, and otherwise the other firm would experiment immediately. Denote the surplus generated by this strategy by S_1^P.

2. *Simultaneous experimentation*: this would involve having both firms experiment immediately at $t = 0$. Denote the surplus generated by this strategy by S_2^P.

It is clear that no other strategy could be optimal for the planner. Moreover, S_1^P and S_2^P have simple expressions. In particular,

$$S_1^P = 2p\Pi_2 + (1 - p)\,p\Pi_1. \tag{3.1}$$

Intuitively, one of the firms experiments first and is successful with probability p. When this happens, the other firm copies a successful innovation, with total payoff $2\Pi_2$. With the complementary probability, $1 - p$, the first firm is unsuccessful, and the second firm experiments independently, with expected payoff $p\Pi_1$. These payoffs occur immediately after the first experimentation and thus are not discounted.

The alternative is to have both firms experiment immediately, which generates expected surplus

$$S_2^P = 2p\Pi_1. \tag{3.2}$$

The comparison of S_1^P and S_2^P implies that simultaneous experimentation by both firms is optimal when $2\beta < 1+p$. In contrast, when $2\beta > 1+p$, the optimal allocation involves one of the firms experimenting first, and the second firm copying successful innovations. This is stated in the next proposition (proof in the text).

Proposition 6. *Suppose that*

$$2\beta \geq 1+p, \tag{3.3}$$

then the optimal allocation involves staggered experimentation, that is, experimentation by one firm and copying of successful innovations. If (3.3) does not hold, then the optimal allocation involves immediate experimentation by both firms. When $2\beta = 1+p$, both staggered experimentation and immediate experimentation are socially optimal.

Note at this point that if consumer surpluses from a new innovation and from the two firms implementing the same project were, respectively, C_1 and $2C_2$, then we would have

$$\begin{aligned} S_1^P &= 2p\left(\Pi_2 + C_2\right) + (1-p)\,p\left(\Pi_1 + C_1\right) \\ S_2^P &= 2p\left(\Pi_1 + C_1\right). \end{aligned}$$

Denoting $\gamma \equiv \left(\Pi_2 + C_2\right) / \left(\Pi_1 + C_1\right)$, it is then clear that condition (3.3) would be replaced by $2\gamma \geq 1+p$ and the rest of the analysis would remain unchanged. If $C_2 = \kappa\Pi_2$ and $C_1 = \kappa\Pi_1$, then this condition would be identical to (3.3).

Let us now compare this to the equilibria characterized so far. Clearly, asymmetric equilibria are identical to the first strategy of the planner and thus generate surplus S_1^P (recall subsection 2.1.). In contrast, the (unique) symmetric equilibrium generates social surplus

$$\begin{aligned} S^E &= \int_0^\infty 2\lambda^* e^{-(2\lambda^*+r)t}\left[2p\Pi_2 + (1-p)\,p\Pi_1\right]dt \\ &= \frac{2\lambda^*}{2\lambda^*+r}\left[2p\Pi_2 + (1-p)\,p\Pi_1\right] = 2p\Pi_2, \end{aligned} \tag{3.4}$$

where λ^* is the (constant) equilibrium flow rate of experimentation given by (2.5). The first line of (3.4) applies because the time of first experimentation corresponds to the first realization of one of two random variables, both with an exponential distribution with parameter λ^* and time is discounted at the rate r. If the first experimentation is successful, which has probability p, surplus is equal to $2\Pi_2$, and otherwise (with probability $1-p$), the second firm also experiments, with expected payoff $p\Pi_1$. The second line is obtained by solving the integral and substituting for (3.1). An alternative way to obtain that $S^E = 2p\Pi_2$ is by noting that at equilibrium the two firms are mixing with a constant flow of experimentation for all times, thus the expected payoff for each should be equal to the payoff when they experiment at time $t = 0$, i.e., $p\Pi_2$.

A straightforward comparison shows that S^E is always (strictly) less than both S_1^P and S_2^P. Therefore, the unique symmetric equilibrium is always inefficient. Moreover, this inefficiency can be quantified in a simple manner. Let $S^P = \max\{S_1^P, S_2^P\}$ and consider the ratio of equilibrium social surplus to the social surplus in the optimal allocation as a measure of inefficiency:

$$\mathfrak{s} \equiv \frac{S^E}{S^P}.$$

Naturally, the lower is $\mathfrak{s}$ the more inefficient is the equilibrium.

Clearly, $\mathfrak{s} < 1$, so that the equilibrium is always inefficient as stated above. More specifically, let us first suppose that (3.3) holds. Then, the source of inefficiency is delayed experimentation. In this case,

$$\begin{aligned} \mathfrak{s} &= \frac{S^E}{S_1^P} \\ &= \frac{2\lambda^*}{2\lambda^* + r} = \frac{2\beta}{2\beta + 1 - p}, \end{aligned}$$

where the last equality simply uses (2.5). It is clear that $\mathfrak{s}$ is minimized, for given p, as

$\beta = (1+p)/2$ (its lower bound given (3.3)). In that case, we have

$$\mathfrak{s} = \frac{1+p}{2}.$$

In addition, as $p \downarrow 0$, $\mathfrak{s}$ can be as low as $1/2$.

Next consider the case where (3.3) does not hold. Then

$$\begin{aligned} \mathfrak{s} &= \frac{S^E}{S_2^P} \\ &= \frac{2\lambda^*}{2\lambda^* + r} \frac{2p\Pi_2 + (1-p)\, p\Pi_1}{2p\Pi_1} = \beta, \end{aligned}$$

where the last equality again uses (2.5) and the definition of β from (2.1). Since this expression applies when $\beta < 1 + p$, β can be arbitrarily small as long as p is small (to satisfy the constraint that $\beta > p$), and thus in this case $\mathfrak{s} \downarrow 0$. In both cases, the source of inefficiency of the symmetric equilibrium is because it generates insufficient incentives for experimentation. In the first case this exhibits itself as delayed experimentation, and in the second, as lack of experimentation by one of the firms.

This discussion establishes (proof in the text).

Proposition 7. 1. *Asymmetric equilibria are Pareto optimal and maximize social surplus when (3.3) holds, but fail to maximize social surplus when (3.3) does not hold.*

2. *The unique symmetric equilibrium is always Pareto suboptimal and never maximizes social surplus. When (3.3) holds, this equilibrium involves delayed experimentation, and when (3.3) does not hold, there is insufficient experimentation.*

3. *When (3.3) holds, the relative surplus in the equilibrium compared to the surplus in the optimal allocation, $\mathfrak{s}$, can be as small as 1/2. When (3.3) does not hold, the symmetric equilibrium can be arbitrarily inefficient. In particular, $\mathfrak{s} \downarrow 0$ as $p \downarrow 0$ and $\beta \downarrow 0$.*

It is straightforward to verify that the results in this proposition apply exactly if consumer surpluses are proportional to firm profits, i.e., $C_2 = \kappa \Pi_2$ and $C_1 = \kappa \Pi_1$. If this is not the case, then the worst-case scenario considered in part 3 can become even worse because of the misalignment between firm profits and consumer surplus resulting from different types of successful research projects.

3.2 Patents

The previous section established the inefficiency of the symmetric equilibrium resulting from delayed and insufficient experimentation. In this section, we discuss how patents can solve or ameliorate this problem. Our main argument is that a patent system provides incentives for greater experimentation or for experimentation without delay.

We model a simple patent system, whereby a patent is granted to any firm that undertakes a successful innovation. If a firm copies a patented innovation, it has to make a payment (compulsory license fee) η to the holder of the patent. We discuss the relationship between this payment and licensing fees in the next section. An appropriately-designed patent system (i.e., the appropriate level of η) can achieve two objectives simultaneously. First, it can allow firms to copy others when it is socially beneficial for the knowledge created by innovations to spread to others (and prevent it when it is not beneficial). Second, it can provide compensation to innovators, so that incentives to free-ride on others are weakened. In particular, when staggered experimentation is optimal, a patent system can simultaneously provide incentives to one firm to innovate early and to the other firm to copy an existing innovation. When η is chosen appropriately, the patent system provides incentives for the *ex post* transfer of knowledge. However, more crucially, it also encourages innovation because an innovation that is copied becomes more profitable than copying another innovation and paying the patent fee. The key here is that the incentives provided by the patent system are "conditional" on whether the other firm has experimented or not, and thus induce an "asymmetric" response from the two firms. This makes innovation

relatively more profitable when the other firm copies and less profitable when the other firm innovates. This incentive structure encourages one of the firms to be the innovator precisely when the other firm is copying. Consequently, the resulting equilibria resemble asymmetric equilibria. Moreover, these asymmetric incentives imply that, when the patent system is designed appropriately, a symmetric equilibrium no longer exists. It is less profitable for a firm to innovate when the other firm is also innovating, because innovation no longer brings patent revenues. Conversely, it is not profitable for a firm to wait when the other firm waits, because there is no innovation to copy in that case.

Our main result in this section formalizes these ideas. We state this result in the following proposition and then provide most of the proof, which is intuitive, in the text.

Proposition 8. *Consider the model with two firms. Suppose that Assumption 1 holds. Then:*

1. *When (3.3) holds, a patent system with*

$$\eta \in \left[\frac{(1-p)\,\Pi_1}{2}, \Pi_2 - p\Pi_1\right)$$

(which is feasible in view of (3.3)), implements the optimal allocation, which involves staggered experimentation, in all equilibria. That is, in all equilibria one firm experiments first, and the other one copies a successful innovation and experiments immediately following an unsuccessful experimentation.

2. *When (3.3) does not hold, then the optimal allocation, which involves simultaneous experimentation, is implemented as the unique equilibrium by a patent system with*

$$\eta > \Pi_2 - p\Pi_1.$$

That is, there exists a unique equilibrium in which both firms immediately experiment.

Let us start with the first claim in Proposition 8. We outline the argument for why $\eta < \Pi_2 - p\Pi_1$ implies that there exists an equilibrium with staggered experimentation, and $\eta \geq \frac{(1-p)\Pi_1}{2}$ ensures that other equilibria, which involve delayed experimentation, are ruled out. Observe that since $\eta < \Pi_2 - p\Pi_1$, the equilibrium involves copying of a successful innovation by a firm that has not acted yet. However, incentives for delaying to copy are weaker because copying now has an additional cost η, and innovation has an additional benefit η if the other firm is imitating. Suppose that firm $\sim i$ will innovate at some date $T > 0$ (provided that firm i has not done so until then). Then the payoffs to firm i when it chooses experimentation and waiting are

$$\begin{aligned} \text{experiment now} &= p\left(\Pi_2 + \eta\right) \\ \text{wait} &= e^{-rT}\left(p\left(\Pi_2 - \eta\right) + (1-p)\, p\Pi_1\right). \end{aligned}$$

It is clear that for any $T > 0$, experimenting is a strict best response, since

$$p\left(\Pi_2 + \eta\right) \geq p\left(\Pi_2 - \eta\right) + (1-p)\, p\Pi_1$$

given that $\eta \geq \frac{(1-p)\Pi_1}{2}$. So experimenting immediately against a firm that is waiting is optimal. To show that all equilibria implement the optimal allocation, we also need to show that both firms experimenting immediately is not an equilibrium. Suppose they did so. Then the payoff to each firm, as a function of whether they experiment or wait, would be

$$\begin{aligned} \text{experiment now} &= p\Pi_1 \\ \text{wait} &= p\left(\Pi_2 - \eta\right) + (1-p)\, p\Pi_1. \end{aligned}$$

Waiting is a strict best response since

$$p(\Pi_2 - \eta) + (1-p)\, p\Pi_1 > p\Pi_1$$

which holds in view of the fact that $\eta < \Pi_2 - p\Pi_1$. This argument makes it intuitive that patents induce an equilibrium structure without delay: waiting is (strictly) optimal when the other firm is experimenting immediately and experimenting immediately is (strictly) optimal when the other firm is waiting. To establish this claim formally, we need to prove that there are no mixed strategy equilibria. This is done in the next lemma.

Lemma 2. *When equation (3.3) holds, there does not exist any equilibrium with mixing.*

Proof. Let us write the expected present discounted value of experimenting at time t for firm i when firm $\sim i$ experiments at the flow rate $\lambda(t)$ as in (2.7) in the proof of Proposition 3 except that we now take patent payments into account and use equation (3.3) so that copying a successful innovation is still profitable. This expression is

$$V(t) = \int_0^t \lambda(z)\, e^{-m(z)} e^{-rz} \left[p(\Pi_2 - \eta) + (1-p)\, p\Pi_1\right] dz + e^{-m(t)} e^{-rt} p(\Pi_2 + \eta),$$

where $m(t)$ is given by (2.6) in the proof of Proposition 3. This expression must be constant for all t in the support of the mixed-strategy equilibrium. The argument in the proof of Proposition 3 establishes that $\lambda(t)$ must satisfy

$$\lambda(t)\left[p(\Pi_2 - \eta) + (1-p)\, p\Pi_1\right] = (r + \lambda(t))\, p(\Pi_2 + \eta).$$

It can be verified easily that since $\eta \geq \frac{(1-p)\Pi_1}{2}$ this equation cannot be satisfied for any $\lambda(t) \in \mathbb{R}_+$ (for any t). Therefore, there does not exist any equilibrium with mixing. ■

Let us next turn to the second claim in the proposition. Suppose that (3.3) is not satisfied and let $\eta > \Pi_2 - p\Pi_1$. Then it is not profitable for a firm to copy a successful

innovation. Therefore, both firms have a unique optimal strategy which is to experiment immediately, which coincides with the optimal allocation characterized in Proposition 6.

The intuition for the results in Proposition 8 can also be obtained by noting that the patent system is inducing experimenters to internalize the externalities that they create. Let us focus on part 1 and suppose that firm 1 experiments while firm 2 delays and copies a successful innovation by firm 1. In this case, the social surplus is equal to $2p\Pi_2 + (1-p)\,p\Pi_1$. Firm 1 only receives $p\Pi_2$ without a patent, and if it were to deviate and delay experimentation, firm 2 would instead receive $p\Pi_2$. Thus to internalize the positive externality that it is creating, firm 1 needs to be compensated for $(1-p)\,p\Pi_1$. A licensee fee of $\eta \geq \frac{(1-p)\Pi_1}{2}$ achieves this, since by experimenting firm 1 receives this license fee with probability p and by delaying, it would have had to pay the same license fee with probability p (and thus $2p\eta \geq (1-p)\,p\Pi_1$). The requirement that $\eta < \Pi_2 - p\Pi_1$ then simply ensures that firm 2 indeed wishes to copy the innovation despite the license fee.

The preceding discussion and Proposition 8 show how an appropriately-designed patent system can be useful by providing stronger incentives for experimentation. When simultaneous experimentation by all parties is socially beneficial, a patent system can easily achieve this by making copying (or "free-riding") unprofitable. On the other hand, when ex post transfer of knowledge is socially beneficial, the patent system can instead ensure this while also preventing delays in all equilibria. It is important to emphasize that, in the latter case, the patent system provides such incentives *selectively*, so that only one of the firms engages in experimentation and the other firm potentially benefits from the innovation of the first firm. In contrast to patents, simple subsidies to research could not achieve this objective. This is stated in the next proposition and highlights the particular utility of a patent system in this environment.

Proposition 9. *Suppose that equation (3.3) holds. Consider a direct subsidy $w > 0$ given to a firm that experiments. There exists no $w \geq 0$ such that all equilibria with subsidies correspond to the optimal allocation.*

Proof. This is straightforward to see. If $w \geq \Pi_2 - p\Pi_1$, there exists an equilibrium in which both firms experiment immediately and if $w < \Pi_2 - p\Pi_1$, the symmetric mixed-strategy equilibrium with delayed experimentation survives. ∎

It is also clear that the same argument applies to subsidies to successful innovation or any combination of subsidies to innovation and experimentation.

3.3 Patents and License Fees

The analysis in the previous section assumed that a firm can copy a successful innovation and in return it has to make some pre-specified payment (compulsory license fee) η to the original innovator. In practice patents often provide exclusive rights to the innovator, who is then allowed to license its product or discovery to other firms. If so, the license fee η would need to be negotiated between the innovator and the (potential) copying firm rather than determined in advance. While such voluntary licensing is an important aspect of the patent system in practice, it is not essential for the theoretical insights we would like to emphasize.

To illustrate this, let us suppose that the copying firm is developing a different but highly substitutable product to the first innovation. Suppose further that the patent system gives exclusive rights to the innovator but if the second firm copies a successful innovation, the court system needs to determine damages. How the court system functions is also part of the patent system. In particular, suppose that if a firm copies a successful innovation without licensing and the innovator brings a lawsuit, it will succeed with probability $\rho \in (0,1)$ and the innovator will receive damages equal to $\kappa(\Pi_1 - \Pi_2)$, where $\kappa > 0$. We ignore legal fees. Given this legal environment, let us interpret η as a license fee negotiated between the potential copying firm and the innovator. For simplicity, suppose that this negotiation can be represented by a take-it-or-leave-it offer by the innovator (this has no effect on the conclusions of this section). If the two firms agree to licensing, their joint surplus is $2\Pi_2$. If they disagree, then the outside option of the copying firm is $\max\{p\Pi_1; \Pi_2 - \rho\kappa(\Pi_1 - \Pi_2)\}$,

where the max operator takes care of the fact that the best alternative for the "copying" firm may be to experiment if there is no explicit licensing agreement. Without licensing, the innovator will receive an expected return of $\Pi_2 + \rho\kappa\left(\Pi_1 - \Pi_2\right)$ if $\Pi_2 - \rho\kappa\left(\Pi_1 - \Pi_2\right) \geq p\Pi_1$ and Π_1 otherwise. This implies that the negotiated license fee, as a function of the parameters of the legal system, will be

$$\eta\left(\rho, \kappa\right) = \begin{cases} \rho\kappa\left(\Pi_1 - \Pi_2\right) & \text{if } p\Pi_1 < \Pi_2 - \rho\kappa\left(\Pi_1 - \Pi_2\right), \\ \Pi_2 - p\Pi_1 & \text{if } p\Pi_1 \geq \Pi_2 - \rho\kappa\left(\Pi_1 - \Pi_2\right) \text{ and } 2\Pi_2 > \Pi_1, \\ \infty & \text{otherwise,} \end{cases}$$

where ∞ denotes a prohibitively expensive license fee, such that no copying takes place. Clearly, by choosing ρ and κ, it can be ensured that $\eta\left(\rho, \kappa\right)$ is greater than $\Pi_2 - p\Pi_1$ when (3.3) does not hold and is between $\left(1 - p\right)\Pi_1/2$ and $\Pi_2 - p\Pi_1$ when it holds. This illustrates how an appropriately-designed legal enforcement system can ensure that equilibrium license fees play exactly the same role as the pre-specified patent fees did in Proposition 8.

Chapter 4

Firms with Heterogeneous Information

Fort the remainder, we relax the assumption that all firms receive signals with identical precision. Instead, now signal quality differs across firms. We continue to assume that each firm receives a positive signal about a single project. But the information content of these signals differs. We parameterize signal quality by the probability with which the indicated project is successful and denote it by p (or by p_i for firm i). We distinguish two cases. First, we discuss the case when signals are publicly known (we limit the discussion to two firms) and, then, we study the case when the signals are drawn from a known distribution represented by the cumulative distribution function $G(p)$ (G is assumed to have strictly positive and continuous density $g(p)$ over its support $[a, b] \subset [0, 1]$) and the realization of p for each firm is independent of the realizations for others and is *private information.* For the case of private signals we discuss the case of two firms in the main text and relegate a discussion on the extension to multiple firms to Appendix B. Finally, throughout we focus on the equivalent of *symmetric equilibria* where strategies do not depend on firm identity.

4.1 Publicly Known Signals

Let p_1, p_2 denote the signals of firms 1 and 2 respectively. We also impose:

Assumption 3.

$$\pi_1 > \pi_2 > \min\{p_1, p_2\} \cdot \pi_1.$$

Note that Assumption 3 implies that firm $i = \arg\min\{p_1, p_2\}$ would find it optimal to copy firm $\sim i$, if the latter was successful at experimentation. Also, note that when $\min\{p_1, p_2\} \cdot \pi_1 \geq \pi_2$ the structure of the equilibrium is straightforward. Let us consider the following two cases: (1) $\max\{p_1, p_2\} \cdot \pi_1 \geq \pi_2$ and (2) $\max\{p_1, p_2\} \cdot \pi_1 < \pi_2$. The next proposition characterizes the unique equilibrium (in fully mixed strategies prior to any experimentation) in both cases (the proof is omitted as it uses similar arguments to that of Proposition 3). For the remainder of the section, let $p_{\max} \equiv \max\{p_1, p_2\}$, $p_{\min} \equiv \min\{p_1, p_2\}$ and similarly $i_{\max} \equiv \arg\max\{p_1, p_2\}$ and $i_{\min} \equiv \arg\min\{p_1, p_2\}$.

Proposition 10. *Suppose that Assumption 3 holds. Then, there exists a unique equilibrium in fully mixed strategies prior to any experimentation. In particular:*

(1) Suppose $p_{\max} \cdot \pi_1 < \pi_2$. Then, in the unique fully mixed equilibrium, firm 1 uses the constant flow rate of experimentation $\lambda_1 = \frac{r \cdot p_2 \cdot \beta}{(1-p_1)p_2 + (p_1 - p_2)\beta}$ and firm 2 uses the rate $\lambda_2 = \frac{r \cdot p_1 \cdot \beta}{(1-p_2)p_1 + (p_2 - p_1)\beta}$. Firm i immediately copies a successful innovation by firm $\sim i$ and experiments if $\sim i$ experiments unsuccessfully.

(2) Suppose $\max\{p_1, p_2\} \cdot \pi_1 \geq \pi_2$. Then, in the unique fully mixed equilibrium, firm $i_{\min}$ uses the constant flow rate of experimentation $\lambda_{\min} = \frac{r\beta}{1-\beta}$ and firm $i_{\max}$ uses the rate $\lambda_{\max} = \frac{r \cdot p_{\min}}{(\beta - p_{\min})p_{\max}}$. Firm $i_{\min}$ immediately copies a successful innovation by firm $i_{\max}$ and experiments if $i_{\max}$ experiments unsuccessfully. On the other hand, if $i_{\min}$ experiments first, then $i_{\max}$ experiments with its own research project (does not copy the potential innovation).

It is worth noting that when $\max\{p_1, p_2\} \cdot \pi_1 \geq \pi_2$, firm $i_{\max}$ delays experimentation not

to copy a potential innovation by firm $i_{\min}$ but so as not to get copied by $i_{\min}$. Proposition 11 is analogous to Proposition 6 and describes the optimal allocation in this setting (proof is omitted).

Proposition 11. *Suppose that*

$$2\beta \geq 1 + p_{\min}, \tag{4.1}$$

then the optimal allocation involves staggered experimentation, that is, experimentation by firm $i_{\max}$ first and copying of successful innovations. If (4.1) does not hold, then the optimal allocation involves immediate experimentation by both firms. When $2\beta = 1 + p_{\min}$, both staggered experimentation and immediate experimentation are socially optimal.

Moreover, we can show that a patent system with:

$$\eta \in \left[\min\left\{\Pi_2 - p_{\max}\Pi_1, \frac{(1 - p_{\min})\, p_{\max}\Pi_1 - (p_{\max} - p_{\min})\,\Pi_2}{p_1 + p_2}\right\}, \Pi_2 - p_{\min}\Pi_1\right)$$

implements the optimal allocation in all equilibria, when (4.1) holds. When (4.1) does not hold, then the optimal allocation is implemented as the unique equilibrium by a patent system with $\eta > \Pi_2 - p_{\min}\Pi_1$ (the claim follows by similar arguments to those in Proposition 8). An interesting feature of the optimal allocation illustrated by Proposition 11 is that it involves a *monotonicity*, whereby the firm with the strongest signal experiments earlier (no later) than the firm with the weaker signal. Yet, this monotonicity does not necessarily hold at equilibrium, since there is a positive probability that the firm with the weaker signal ($i_{\min}$) experiments before the firm with the stronger signal ($i_{\max}$).

4.2 Private Signals

For the remainder, we assume that p's are drawn independently from a known distribution with cumulative distribution function $G(p)$ and continuous density $g(p)$ over its support. As the title of the section indicates, the realization of p's are private information. We start with the following lemma, which follows from the definition of β in (2.1). It will

play an important role in the analysis that follows (proof omitted).

Lemma 3. *Suppose that firm $\sim i$ has innovated successfully. If $p_i > \beta$, firm i prefers to experiment with its own project. If $p_i < \beta$, firm i prefers to copy a successful project.*

Proposition 12 below provides a characterization of the unique symmetric equilibrium with two firms. We show that the equilibrium takes the following form: firms with strong signals (in particular, $p \geq \beta$) experiment immediately, while those with weaker signals (i.e., $p < \beta$) experiment at time $\tau(p)$ with $\tau(\beta) = 0$, unless there has been experimentation at any earlier time. Function $\tau(p)$ is strictly decreasing and maps signals to time of experimentation provided that the other player has not yet experimented. The proof of the proposition uses a series of lemmas and is relegated to the Appendix.

Proposition 12. *Let the support of G be $[a, b] \subset [0, 1]$ and define $\bar{b} \equiv \min\{\beta, b\}$ and*

$$\bar{\tau}(p) \equiv \frac{1}{r\beta G\left(\bar{b}\right)}\left[\log G(\bar{b})\left(1-\bar{b}\right) - \log G(p)\left(1-p\right) + \int_p^{\bar{b}} \log G(z)dz\right]. \tag{4.2}$$

Then the unique symmetric equilibrium involves:

$$\tau\left(p\right) = \begin{cases} 0 & \text{if } p \geq \beta \\ \bar{\tau}\left(p\right) & \text{if } p \in [a, \beta) \end{cases}.$$

That is, firms with $p \geq \beta$ experiment immediately and firms with $p \in [a, \beta)$ experiment at time $\bar{\tau}\left(p\right)$ unless there has been an experimentation at $t < \bar{\tau}\left(p\right)$. If there is experimentation at $t < \bar{\tau}\left(p\right)$, then a firm with $p \in [a, \beta)$ copies it if the previous attempt was successful and experiments immediately if it was unsuccessful.

A particularly simple example to illustrate Proposition 12 is obtained when G is uniform over $[a, b]$ for $0 < a < b \leq \beta$. In that case

$$\tau\left(p\right) = \frac{1}{r\beta}\left[p - \log p - b + \log b\right] \text{ for all } p \in [a, b]. \tag{4.3}$$

An interesting feature of the symmetric equilibria in this case is evident from (4.3): for a arbitrarily close to 0, experimentation may be delayed for arbitrarily long time. It can be verified from (4.2) that this is a general feature (for types arbitrarily close to to the lower support a, $-\log G(p)$ is arbitrarily large).

4.3 Welfare

In this section, we discuss welfare in the environment with private, heterogeneous signals studied in the previous section. In particular, consider a social planner that is interested in maximizing total surplus (as in Section 3.1. What the social planner can achieve will depend on her information and on the set of instruments that she has access to. For example, if the social planner observes the signal quality, p, for each firm, then she can achieve a much better allocation than the equilibrium characterized above. However, it is more plausible to limit the social planner to the same information structure. In that case, the social planner will have to choose either the same equilibrium allocation as in the symmetric equilibria characterized in the previous two sections, or she will implement an asymmetric equilibrium, where one of the firms is instructed to experiment first regardless of its p (this cannot be conditioned on p since p is private information).[1]

More specifically, let us focus on the economy with two firms and suppose that the support of G is $[a, b] \subset [0, \beta]$. In this case, without eliciting information about the realization of firm types, the p's, the planner has three strategies.

1. *Staggered asymmetric experimentation*: in this case, the social planner would instruct one of the firms to experiment immediately and then have the other firm copy if there is a successful innovation. Since the social planner does not know the p's, she has to pick the experimenting firm randomly. We denote the social surplus generated by this strategy by S_1^P.

[1]Yet another alternative is to specify exactly the instruments available to the planner and characterize the solution to a mechanism design problem by the planner. However, if these instruments allow messages and include payments conditional on messages, the planner can easily elicit the necessary information from the firms.

2. *Staggered equilibrium experimentation*: alternatively, the social planner could let the firms play the symmetric equilibrium of the previous two sections, whereby a firm of type p will experiment at time $\tau(p)$ unless there has previously been an experimentation by the other firm. We denote the social surplus generated by this strategy by S^E, since this is the same as the equilibrium outcome.[2]

3. *Simultaneous experimentation*: in this case, the social planner would instruct both firms to experiment immediately. We denote the social surplus generated by this strategy by S_2^P.

The social surpluses from these different strategies are given as follows. In the case of staggered asymmetric experimentation, we have

$$S_1^P = \int_a^b \left[p_1 2\Pi_2 + (1 - p_1) \left(\int_a^b p_2 dG(p_2) \right) \Pi_1 \right] dG(p_1).$$

In contrast, the expected surplus from the unique (mixed-strategy) symmetric equilibrium can be written as

$$S^E =$$

$$\int_a^b e^{-r\tau(\max\{p_1,p_2\})} \left[\max\{p_1, p_2\} 2\Pi_2 + (1 - \max\{p_1, p_2\}) \int_a^b \min\{p_1, p_2\} \Pi_1 \right] dG(p_1) dG(p_2).$$

Intuitively, this expression follows by observing that in the equilibrium as specified in Proposition 20, the firm with the stronger signal (higher p) will experiment first, so there will be delay until $\max\{p_1, p_2\}$. At that point, this firm will succeed with probability $\max\{p_1, p_2\}$, in which case the second firm will copy. If the first firm fails (probability $1-$

[2]Without eliciting information about firm types and without using additional instruments, the social planner cannot implement another monotone staggered experimentation allocation. For example, she could announce that if there is no innovation until some time $t > 0$, one of the firms will be randomly forced to experiment. But such schemes will not preserve monotonicity, since at time t, it may be the firm with lower p that may be picked for experimentation. In the next section, we discuss how she can implement better allocations using patent payments.

$\max\{p_1, p_2\}$), then the second firm experiments and succeeds with probability $\min\{p_1, p_2\}$. Since both p_1 and p_2 are randomly drawn independently from G, we integrate over G twice to find the expected surplus.

The surplus from simultaneous experimentation, on the other hand, takes a simple form and is given by

$$S_2^P = 2\Pi_1 \int_a^b p dG(p),$$

since in this case each firm is successful and generates payoff Π_1 with probability p distributed with distribution function G.

In this case, there is no longer any guarantee that $\max\{S_1^P, S_2^P\} > S^E$. Therefore, the symmetric equilibrium may generate a higher expected surplus (relative to allocations in which the social planner does not have additional instruments). To illustrate this, let us consider a specific example, where p has a uniform distribution over $[0, \beta]$. In this case, staggered asymmetric experimentation gives

$$S_1^P = \int_0^\beta \int_0^\beta 2p_1\Pi_2 + (1-p_1)p_2\Pi_1 dp_2 dp_1 = \Pi_2\left(\frac{1}{2} + \frac{3}{4}\beta\right),$$

whereas simultaneous experimentation gives

$$S_2^P = \int_0^\beta 2p\Pi_1 dp = \beta\Pi_1 = \Pi_2.$$

Comparing simultaneous experimentation and staggered asymmetric experimentation, we can conclude that $S_1^P > S_2^P$ whenever $\beta > 2/3$ and $S_1^P < S_2^P$ whenever $\beta < 2/3$, showing that, as in the case with common signals, either simultaneous or staggered experimentation might be optimal. Next, we can also compare these surpluses to S^E. Since p is uniformly distributed in $[0, \beta]$, (4.2) implies that

$$\tau(p) = \frac{1}{r\beta}\left[p - \log p - \beta + \log \beta\right].$$

As a consequence, $\max\{p_1, p_2\}$ has a Beta(2,1) distribution (over $[0, \beta]$) while $\min\{p_1, p_2\}$ is distributed Beta(1,2). Then evaluating the expression for S^E, we find that when $0 \leq \beta \leq 2/3$, $S_2^P > S^E$, so simultaneous experimentation gives the highest social surplus. When $2/3 \leq \beta \leq \beta^* \simeq 0.895$, $S_1^P > S^E$, so that staggered asymmetric experimentation gives the highest social surplus. Finally, when $\beta^* \leq \beta \leq 1$, $S^E > S_1^P > S_2^P$, so the symmetric equilibrium gives higher social surplus than both staggered asymmetric experimentation and simultaneous experimentation.

Finally, it is also straightforward to see that by choosing G to be highly concentrated around a particular value $\bar{p}$, we show that the symmetric equilibrium can be arbitrarily inefficient relative to the optimal allocation.

4.4 Equilibrium with Patents

Equilibria with patents are also richer in the presence of heterogeneity. Let us again focus on the case in which there are two firms. Suppose that there is a patent system identical to the one discussed in Section 3.2, whereby a firm that copies a successful innovation pays η to the innovator. Let us define

$$p^\eta \equiv \frac{\Pi_2 - \eta}{\Pi_1}. \tag{4.4}$$

It is clear, with a reasoning similar to Lemma 3, that only firms with $p < p^\eta$ will copy when the patent system specifies a payment of η. The next proposition characterizes the structure of equilibria with patents (the proof is relegated to the Appendix).

Proposition 13. *Suppose that there are two firms and the patent system specifies a payment $\eta > 0$ for copying. Let p^η be given by (4.4), the support of G be $[a, b] \subset [0, 1]$, and define*

$\bar{b} \equiv \min\{b, p^\eta\}$ *and*

$$\bar{\tau}^\eta(p) \equiv$$

$$\frac{1}{r(\Pi_2+\eta)G(\bar{b})}\left[\log G(\bar{b})\left(\Pi_1-2\eta-\bar{b}\Pi_1\right)-\log G(p)\left(\Pi_1-2\eta-p\Pi_1\right)+\Pi_1\int_p^{\bar{b}}\log G(z)\,dz\right]. \tag{4.5}$$

Then the unique symmetric equilibrium involves:

$$\tau^\eta(p) = \begin{cases} 0 & \text{if } p \geq p^\eta \\ \bar{\tau}^\eta(p) & \text{if } p \in [a, p^\eta) \end{cases}.$$

That is, firms with $p \geq p^\eta$ experiment immediately and firms with $p \in [a, p^\eta)$ experiment at time $\bar{\tau}^\eta(p)$ unless there has been an experimentation at $t < \bar{\tau}^\eta(p)$.

Moreover, a higher η tends to reduce delay. In particular:

- *for any $\eta' > \eta$ such that $b < p^\eta$ and $b < p^{\eta'}$, we have $\tau^{\eta'}(p) \leq \tau^\eta(p)$ for all $p \in [a, b]$, with strict inequality whenever $\tau^\eta(p) > 0$;*
- *for any η' such that $b > p^{\eta'}$, there exists $p^*(\eta') \in [0, p^{\eta'})$ such that $\tau^\eta(p)$ is decreasing in η starting at $\eta = \eta'$ for all $p \in [p^*(\eta'), p^{\eta'}]$, with strict inequality whenever $\tau^{\eta'}(p) > 0$.*

Note that the first bullet point considers the case when all firms would prefer to copy a successful innovation than to experiment on their own (since $b < p^{\eta'} < p^\eta$), whereas the second bullet point considers the case when there is a positive probability that a firm obtains a strong enough signal and prefers to experiment on its own.

The result highlights an important role of patents in experimentation. When η increases, $\tau(p)$ tends to become "steeper" so that there is less delay and thus "time runs faster". In particular, whenever $p^\eta < b$, $\tau(p)$ is reduced by an increase in patent payments. When $p^\eta > b$, this does not necessarily apply for very low p's, but is still true for high p's.

Overall, this result implies that as in the case with common p's, patents tend to increase experimentation incentives and reduce delay. In the limit, when η becomes arbitrarily large, the equilibrium involves simultaneous experimentation. Nevertheless, as discussed in Section 4.3, simultaneous experimentation may not be optimal in this case.

Alternatively (and differently from Proposition 13), a patent system can also be chosen such that the socially beneficial ex post transfer of knowledge takes place. In particular, suppose that there has been an innovation and the second firm has probability of success equal to p. In this case, social surplus is equal to $2\Pi_2$ if there is copying, and it is equal to $\Pi_1 + p\Pi_1$ if the second firm is forced to experiment. This implies that to maximize ex post social welfare, firms with $p \leq 2\beta - 1$ should be allowed to copy, whereas firms with $p > 2\beta - 1$ should be induced to experiment. Clearly, from (4.4) choosing $\eta = \Pi_1 - \Pi_2$ achieves this. Naturally, from Proposition 13, this will typically lead to an equilibrium with staggered experimentation. This argument establishes the following proposition (proof in the text).

Proposition 14. *A patent system with $\eta = \Pi_1 - \Pi_2$ induces the socially efficient copying and experimentation behavior for all $p \in [a, b]$, but typically induces delayed experimentation.*

The juxtaposition of Propositions 13 and 14 implies that when signal quality is heterogeneous and private information, the patent system can ensure either rapid experimentation or the socially beneficial ex post transfer of knowledge (and experimentation by the right types), but will not typically be able to achieve both objectives simultaneously. However, appropriately designed patents typically improve efficiency, as is stated in the following corollary. Moreover, note that unless the types distribution, i.e., G, is skewed towards low signals, the optimal patent payment will satisfy $\eta^* \geq \Pi_1 - \Pi_2$.

Corollary 1. *Suppose that there are two firms and the patent system specifies a payment $\eta > 0$ for copying. Then, the aggregate payoff of the firms is higher than the case when*

$\eta = 0$, unless both firms have very weak signals, i.e., $p_1, p_2 \le p^*(\eta)$, where $p^*(\eta) < p^\eta$ is a constant.

Chapter 5

A Model of Information Exchange in Social Networks

We start by presenting the model for a finite set $\mathcal{N}^n = \{1, 2, \cdots, n\}$ of *agents*. As we are interested in economies with a large number of agents, we then describe the limit economy as $n \to \infty$.

5.1 Actions, Payoffs and Information

Each agent $i \in \mathcal{N}^n$ chooses an action $x_i \in \Re$. Her payoff depends on her action and an underlying *state of the world* θ, which is an exogenous random variable. In particular, agent i's payoff when she takes action x_i and the state of the world is θ is given by $\pi_i = f(x_i, \theta)$. We impose the following assumption on f:

Assumption 4. *The payoff function $f(x, \theta)$ is quadratic in both x and θ. Moreover, f is twice-differentiable and $f_{xx} < 0$, $f_{x\theta} \neq 0$.*

Assuming that f is quadratic greatly simplifies subsequent analysis. We also assume that $f_{xx} < 0$ (concavity with respect to one's action), so as to guarantee that best responses are well-defined. Finally, $f_{x\theta} \neq 0$ ensures that θ is relevant for the behavior of an individual. To ease exposition, we let $f(x, \theta) = \pi - (x - \theta)^2$, where π is a constant. It is straightforward

to extend our results to any function that satisfies Assumption 4.

The state of the world θ is unknown and agents observe noisy signals about its realization. In particular, θ is drawn from a Normal distribution with known mean μ and variance ρ. Each agent receives a private signal $s_i = \theta + z_i$, where the z_i's are idiosyncratic noises, independent from one another and θ, with common mean $\bar{\mu}$, normalized to 0, and variance $\bar{\rho}$.

5.2 Communication

Our focus is on information aggregation, when agents are embedded in a social network structure that can be thought of as imposing communication constraints to the agents. In particular, agent i forms beliefs about the state of the world from her *private signal* s_i, as well as information she can obtain from other agents through a given *communication network* G^n, which will be described shortly. We assume that time $t \in [0, \infty)$ is continuous and there is a common discount rate $r > 0$. At a given time instant t, agent i decides whether to take an irreversible *action* x_i (and receive payoff $f(x_i, \theta)$ discounted by e^{-rt}) or "wait" so as to obtain more information from her peers. Throughout the rest of the thesis, we say that the agent "exits" at time t, if she chooses to take an irreversible action at time t. Discounting implies that an earlier exit is preferred to a later one. In terms of notation, we let u_i^n denote the utility of agent i, when the number of agents in the society is n. From above, $u_i^n = e^{-rt} f(x_i, \theta)$, when agent i takes action x_i at time t.

Finally, we describe how agent i obtains information from her peers in the social network structure. Let the directed graph $G^n = (\mathcal{N}^n, \mathcal{E}^n)$, where $\mathcal{N}^n = \{1, \cdots, n\}$ is the set of agents and $\mathcal{E}^n$ is the set of directed edges with which agents are linked, represent the communication network, in which agents are embedded. We say that agent j can *obtain information* from i or that agent i can *send information* to j if there is an edge from i to j in graph G^n, i.e., $(i, j) \in \mathcal{E}^n$. Let $I_{i,t}^n$ denote the *information set* of agent i at time t and $\mathcal{I}_{i,t}^n$ denote the set of all possible information sets. Then, for every pair of agents i, j, such

that $(i,j) \in \mathcal{E}^n$, we say that agent j *communicates* with agent i or that agent i sends a *message* to agent j, and define the following mapping

$$m^n_{ij,t} : \mathcal{I}^n_{i,t} \to \mathcal{M}^n_{ij,t} \text{ for } (i,j) \in \mathcal{E}^n,$$

where $\mathcal{M}^n_{ij,t}$ denotes the set of messages that agent i can send to agent j at time t. The definition of $m^n_{ij,t}$ captures the fact that communication is directed and is only allowed between agents that are linked in the communication network, i.e., j communicates with i if and only if $(i,j) \in \mathcal{E}^n$. The direction of communication should be clear: when agent j communicates with agent i, then agent i sends a message to agent j, that could in principle depend on the information set of agent i as well as the identity of agent j. Importantly, we assume that the cardinality ("dimensionality") of $\mathcal{M}^n_{ij,t}$ is no less than that of $\mathcal{I}^n_{i,t}$, so that communication can take the form of agent i sharing all her information with agent j. This has two key implications. First, an agent can communicate (indirectly) with a much larger set of agents than just her immediate neighbors, albeit with a time delay. Second, mechanical duplication of information can be avoided. For example, the second time agent j communicates with agent i, she can repeat her original signal, but this will not be recorded as an additional piece of information by agent j, since given the size of the message space $\mathcal{M}^n_{ij,t}$, each piece of information can be "tagged". This ensures that under truthful communication, there need be no confounding of new information and previously communicated information. Figure 5-1 illustrates the process of information aggregation centered at a particular agent.

The times at which communication takes place are exponentially distributed with parameter $\lambda > 0$.[1] Let T_t denote the set of times that agents communicated with their

[1] Alternatively, agents "wake" up and communicate with their neighbors, when a Poisson clock with rate λ ticks.

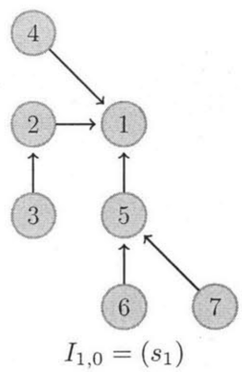

(a) Time $t = 0$.

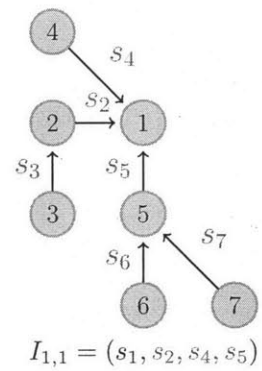

(b) First communication step.

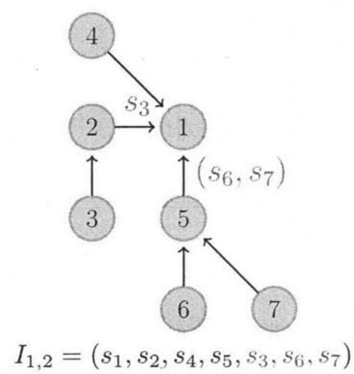

(c) Second communication step.

Figure 5-1: The information set of agent 1 under truthful communication.

neighbors up to time t. That defines the information set of agent i at time $t > 0$ as:

$$I_{i,t}^n = \{s_i, m_{ji,\tau}^n, \text{ for all } \tau \in T_t \text{ and } j \text{ such that } (j,i) \in \mathcal{E}^n\}$$

and $I_{i,0}^n = \{s_i\}$. In particular, the information set of agent i at time $t > 0$ consists of her private signal and all the messages her neighbors sent to i in previous communication times. Agent i's action at time t is a mapping from her information set to the set of actions, i.e.,

$$\sigma_{i,t}^n : \mathcal{I}_{i,t}^n \to \{\text{“wait”}\} \cup \Re.$$

The tradeoff between taking an irreversible action and waiting, should be clear at this point. An agent would wait, in order to communicate with a larger set of agents and choose a better action. On the other hand, future is discounted, therefore, delaying is costly.

We close the section with a number of definitions. We define a *path* between agents i and j in network G^n as a sequence $i_1, \cdots, i_K$ of distinct nodes such that $i_1 = i$, $i_K = j$ and $(i_k, i_{k+1}) \in \mathcal{E}^n$ for $k \in \{1, \cdots, K-1\}$. The *length* of the path is defined as $K - 1$. Moreover, we define the distance of agent i to agent j as the length of the shortest path from i to j in network G^n, i.e.,

$$dist^n(i,j) = \min\{\text{length of } \mathcal{P} \mid \mathcal{P} \text{ is a path from } i \text{ to } j \text{ in } G^n\}.$$

Finally, the (indirect) *neighborhood* of agent i at communication step r is defined as

$$B_{i,r}^n = \{j \mid dist^n(j,i) \leq t\},$$

where $B_{i,0}^n = \{i\}$, i.e., $B_{i,r}^n$ consists of all agents that are at most r *links* away from agent i in graph G^n. Intuitively, if agent i waits for r communication steps and all of the intervening agents receive and communicate information truthfully, i will have access to all of the

signals of the agents in the set $B_{i,r}^n$.

5.3 Equilibria of the Information Exchange Game

We refer to the game defined above as the *Information Exchange Game.* We next define the equilibria of the information exchange game $\Gamma_{info}(G^n)$, given that the communication network is given by G^n. Note that we use the standard notation σ_{-i} to denote the strategies of agents other than i. Also, we let $\sigma_{i,-t}$ denote the vector of actions of agent i at all times except t.

Definition 1. *An action strategy profile $\sigma^{n,*}$ is a pure-strategy Perfect Bayesian Equilibrium of the information exchange game $\Gamma_{info}(G^n)$ if for every $i \in \mathcal{N}^n$ and time t, $\sigma_{i,t}^{n,*}$ maximizes the expected payoff of agent i given the strategies of other agents $\sigma_{-i}^{n,*}$, i.e.,*

$$\sigma_{i,t}^{n,*} \in \arg \max_{y \in \{\text{“}wait\text{”}\} \cup \Re} \mathbb{E}_{\left((y, \sigma_{i,-t}^{n,*}), \sigma_{-i}^{n,*}\right)}(u_i^n \big| I_{i,t}^n).$$

We denote the set of equilibria of this game by $INFO(G^n)$.

For the remainder, we refer to a pure-strategy Perfect Bayesian Equilibrium simply as equilibrium (we do not study mixed strategy equilibria).

Note that if agent i decides to exit and take an action at time t, then her optimal action would be:

$$x_{i,t}^{n,*} = \arg\max_x \mathbb{E}[f(x, \theta) \big| I_{i,t}^n] = \mathbb{E}[\theta \big| I_{i,t}^n],$$

where the second equality holds when $f(x, \theta) = \pi - (x - \theta)^2$. Thus, given that actions are irreversible, the agent's decision problem reduces to determining the timing of her action. It is straightforward to see that at equilibrium an agent takes an irreversible action immediately after some communication step concludes. Thus, an equilibrium strategy profile σ induces an equilibrium timing profile $\tau^{n,\sigma}$, where $\tau_i^{n,\sigma}$ designates the communication step after which agent i exits by taking an irreversible action. The τ notation is convenient to use for the statement of some of our results below.

5.4 Learning in Large Societies

We are interested in whether equilibrium behavior leads to information aggregation. This is captured by the notion of "asymptotic learning", which characterizes the behavior of agents over communication networks with growing size. We first focus on learning over a fixed communication network, i.e., we study agents' decisions along equilibria of the information exchange game.

We consider a sequence of communication networks $\{G^n\}_{n=1}^{\infty}$ where $G^n = \{\mathcal{N}^n, \mathcal{E}^n\}$ with $\mathcal{N}^n = \{1, \cdots, n\}$ and refer to this sequence of communication networks as a *society*. A sequence of communication networks induces a sequence of information exchange games, and with a slight abuse of notation we use the term *equilibrium* to denote a sequence of equilibria of the sequence of information exchange games, or of the society $\{G^n\}_{n=1}^{\infty}$. We denote such an equilibrium by $\sigma = \{\sigma^n\}_{n=1}^{\infty}$, which designates that $\sigma^n \in INFO(G^n)$ for all n. For any fixed $n \geq 1$ and any equilibrium of the information exchange game $\sigma^n \in INFO(G^n)$, we introduce the indicator variable:

$$M_i^{n,\epsilon} = \begin{cases} 1 & \text{if agent } i \text{ takes an action that is } \epsilon\text{-"close" to the optimal,} \\ 0 & \text{otherwise.} \end{cases} \tag{5.1}$$

In other words, $M_i^{n,\epsilon} = 1$ if and only if agent i chooses irreversible action x_i, such that $|x_i - x^*| \leq \epsilon$, where x^* denotes the optimal action for agent i given complete information (i.e., knowing the realization of θ).

Next definition introduces ϵ, δ-asymptotic learning for a given society.

Definition 2. *We say that ϵ, δ-asymptotic learning occurs in society $\{G^n\}_{n=1}^{\infty}$ along equilibrium σ if we have:*

$$\lim_{n \to \infty} \mathbb{P}_\sigma \left(\left[\frac{1}{n} \sum_{i=1}^{n} (1 - M_i^{n,\epsilon}) \right] > \epsilon \right) < \delta.$$

This definition[2] equates ϵ, δ-asymptotic learning occurs with all but an ϵ- fraction of the agents taking an action that is ϵ-"close" to their optimal action (as the society grows infinitely large) with probability at least $1 - \delta$. Moreover, we define perfect asymptotic learning as follows.

Definition 3. *We say that perfect asymptotic learning occurs in society $\{G^n\}_{n=1}^{\infty}$ along equilibrium σ if we have:*

$$\lim_{n \to \infty} \mathbb{P}_\sigma \left(\left[\frac{1}{n} \sum_{i=1}^{n} (1 - M_i^{n,\epsilon}) \right] > \epsilon \right) = 0.$$

for any $\epsilon > 0$.

Perfect asymptotic learning is, thus, a strong definition and requires all but a negligible fraction of the agents taking essentially their optimal action.

5.5 Assumptions on the Information Exchange Process

The communication model described in Section 5.2 is fairly general. In particular, we did not restrict the set of messages that an agent can send or specify her information set. Throughout, we maintain the assumption that the communication network G^n is *common knowledge.* Also, we focus on the following three environments of increasing complexity, defined by Assumptions 5, 6 and 7 respectively.

Assumption 5 (Continuous Communication). *Communication between agents is continuous if*

$$m_{ij,t}^n = \{s_\ell \text{ for all } \ell \in B_{i,r}^n\},$$

for all agents i, j and time periods t, where $r = |T_t|$, i.e., the number of times that commu-

[2]Note that we could generalize Definition 2 by introducing yet another parameter and study ϵ, δ, ζ-asymptotic learning, in which case we would require that $\lim_{n\to\infty} \lim_{t\to\infty} \mathbb{P}_\sigma \left(\left[\frac{1}{n} \sum_{i=1}^{n} (1 - M_i^{n,\epsilon}) \right] > \zeta \right) < \delta$.

nication took place until time t.

This assumption is adopted as a prelude to Assumptions 6 and 7, because it is simpler to work with and as we show the main results that hold under this assumption, generalize to the more complex environments generated by Assumptions 6 and 7. Intuitively, this assumption compactly imposes three crucial features: (1) As already noted, communication takes place by sharing signals, so that when agent j communicates with agent i at time t, then agent i sends to j all the information agent i has obtained thus far, i.e., the private signals of all agents that are at a distance at most $r = |T_t|$ from i (refer back to Figure 5-1 for an illustration of the communication process centered at a particular agent); (2) Communication is *continuous* in the sense that agents do not stop transmitting new information even after taking their irreversible action. This also implies that agents never *exit* the social network, which would be a good approximation to friendship networks that exist for reasons unrelated to communication; (3) Agents cannot strategically manipulate the messages they sent, i.e., an agent's private signal is *hard information.*

Assumption 6 relaxes the second feature above, the continuous transmission of information.

Assumption 6 (Truthful Communication). *Communication between agents is truthful, i.e.,*

$$m_{ij,t}^n = \{s_\ell \text{ for all } \ell \in I_{i,t}^n\},$$

for all agents i, j and time periods t.

Intuitively, it states that when an agent takes an irreversible action, then she no longer obtains new information and, thus, can only communicate the information she has obtained until the time of her decision. The difference between Assumptions 5 and 6 can be seen from the fact that in Assumption 6 we write $I_{i,t}^n$ as opposed to $B_{i,r}^n$, which implies that an agent stops receiving and subsequently communicating new information as soon as she takes an irreversible action. We believe that this assumption is a reasonable approximation

to communication in social networks, since an agent that engages in information exchange to make a decision would have weaker incentives to collect new information after reaching that decision. Nevertheless, she can still communicate the information she had previously obtained to other agents. We call this type of communication *truthful* to stress the fact that the agents cannot strategically manipulate the information they communicate.[3]

Finally, we discuss the implications of relaxing Assumption 6 by allowing *strategic communication*, i.e., when agents can strategically lie or babble about their information. In particular, we replace Assumption 6 with Assumption 7.

Assumption 7 (Strategic Communication). *Communication between agents is strategic if*

$$m_{ij,t}^n \in \Re^{\left|I_{i,t}^n\right|},$$

for all agents i, j and time periods t.

This assumption makes it clear that in this case the messages need not be truthful. Allowing strategic communication adds an extra dimension in an agent's strategy, since the agent can choose to "lie" about (part) of her information set with some probability, in the hope that this increases her expected payoff. Note that, in contrast with "cheap talk" models, externalities in our framework are purely informational as opposed to payoff relevant. Thus, an agent may have an incentive to "lie" as a means to obtain more information from the information exchange process.

[3]Yet another variant of this assumption would be that agents exit the social network after taking an action and stop communicating entirely. In this case, the results are again similar if their action is observed by their neighbors. If they exit the social network, stop communication altogether and their action is not observable, then the implications are different. We do not analyze these variants in the current version to save space.

Chapter 6

Learning and Efficient Communication

In this chapter, we present our main results on learning and discuss their implications for the aggregate welfare of the agents. Before doing so, we discuss the decision problem of a single agent, i.e., determining the best time to take an irreversible action given that the rest of the agents behave according to strategy profile σ. Later, we contrast the single agent problem with that of a social planner, whose objective is to maximize the expected aggregate welfare.

6.1 The single agent problem

Given our normality assumption on both θ and the private signals as well as quadratic preferences, i.e., $f(x,\theta) = \pi - (x-\theta)^2$, the expected payoff of agent i taking an action after observing k private signals (including her own) is given by:

$$\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k},$$

discounted by the appropriate factor. Let $V_i(I_{i,t}^n, \sigma)$ denote agent i's optimal value function at information set $I_{i,t}^n$, when the rest of the agents behave according to strategy profile σ.

Then, by the principle of optimality,

$$V_i(I_{i,t}^n,\sigma) = \max \begin{cases} \pi - \frac{1}{1/\rho^2+1/\bar{\rho}^2 k_{i,t}^{n,\sigma}} & \text{(when she takes the optimal irreversible action)}, \\ e^{-rdt}\ \mathbb{E}[V_i(I_{i,t+dt}^n,\sigma)] & \text{(when she decides to wait, i.e., } x = \text{“}wait\text{”)}, \end{cases}$$

where $k_{i,t}^{n,\sigma}$ denotes the number of distinct private signals agent i has observed up to time t. The first line is equal to the expected payoff for the agent when she chooses the optimal irreversible action under information set $I_{i,t}^n$, i.e., $\mathbb{E}[\theta|I_{i,t}^n]$, and she has observed $k_{i,t}^{n,\sigma}$ private signals, while the second line is equal to the discounted expected continuation payoff. For the latter, we have that with probability λdt, communication takes place in time interval $[t, t+dt]$, thus the information set of agent i expands; with probability $(1-\lambda dt)$ there is no communication and the value function for agent i remains unchanged. If communication takes place in interval $[t, t+dt]$, then agent i observes $|B_{i,|T_t|+1}^{n,\sigma}| - |B_{i,|T_t|}^{n,\sigma}|$ additional signals, where recall that $|T_t|$ denotes the number of communication rounds up to time t and $B_{i,r}^{n,\sigma}$ denotes the set of agents that are distance at most r from i and can be reached by i under profile σ, i.e., there exists a directed path from agent i to each of those agents in G^n, that has length at most r and $\tau_j^{n,\sigma} \geq r - x$, where $x = dist(i,j)$ (which ensures that the information will be transmitted by j).

Note that since we assume that signals are statistically equivalent and independent (as well as truthful), the value function can simply be expressed as a function of the number of distinct signals in $I_{i,t}^n$, $k_{i,t}^{n,\sigma}$ and profile σ. The agent will choose to take an irreversible action and not wait if

$$\begin{aligned} \pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k_{i,t}^{n,\sigma}} &\geq e^{-rdt}\ \mathbb{E}_\sigma[V_i(I_{i,t+dt}^n,\sigma)] \\ &\geq e^{-rdt}(\lambda dt[V_i(k_{i,t}^{n,\sigma} + |B_{i,|T_t|+1}^{n,\sigma}| - |B_{i,|T_t|}^{n,\sigma}|,\sigma) + (1-\lambda dt)V_i(k_{i,t}^{n,\sigma},\sigma))] \end{aligned}$$

Thus, we obtain that the agent should choose not to wait if:

$$V_i(k_{i,t}^{n,\sigma} + |B_{i,|T_t|+1}^{n,\sigma}| - |B_{i,|T_t|}^{n,\sigma}|, \sigma) \leq \frac{r+\lambda}{\lambda}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k_{i,t}^n}\right).$$

This establishes the following lemma, which states that an agent's optimal action takes the form of a threshold rule: an agent decides to take an irreversible action at time t as long as she has observed more that $k_{\sigma,(i,t)}^{n,*}$ private signals.

Lemma 4. *Let Assumption 6 hold. Given communication network G^n and equilibrium $\sigma \in INFO(G^n)$, there exists a sequence of signal thresholds for each agent i, $\{k_{\sigma,(i,t)}^{n,*}\}_{t=0}^{\infty}$, that depend on the current time period, the agent i, the communication network G^n and σ such that agent i maximizes her expected utility at information set $I_{i,t}^n$ by taking action $x_{i,t}^n(I_{i,t}^n)$ defined as*

$$x_{i,t}^n(I_{i,t}^n) = \begin{cases} \mathbb{E}[\theta|I_{i,t}^n], & \text{if } k_{i,t}^{n,\sigma} \geq k_{\sigma,(i,t)}^{n,*}, \\ \text{“wait”}, & \text{otherwise,} \end{cases}$$

A consequence of Lemma 4 is that an equilibrium strategy profile σ defines both a time in which agent i acts (immediately after communication step $\tau_i^{n,\sigma}$), but also the number of signals that agent i has access to when she acts.

6.2 Asymptotic Learning

We begin the discussion by introducing the concepts that are instrumental for asymptotic learning: *the observation radius* and *k-radius sets.* Recall that an equilibrium of the information exchange game on communication network G^n, $\sigma^n \in INFO(G^n)$, induces a timing profile $\tau^{n,\sigma}$, such that agent i takes an irreversible action after $\tau_i^{n,\sigma}$ communication steps. We call $\tau_i^{n,\sigma}$ the *observation radius* of agent i under equilibrium profile σ^n. Note that under Assumption 5, there is no dependence of an agent's expected payoff on the particular equilibrium strategies, that players follow. Therefore, when Assumption 5 holds, the observation radius of agent i is the same irrespective of the strategy profile followed by the

rest of the agents, thus we may drop superscript σ. We refer to agent i's observation radius τ_i^n under Assumption 5 as the *perfect observation radius.* In the appendix we present an efficient algorithm based on dynamic programming for computing the observation radius of each agent i given that the rest of the agents follow strategy profile σ_{-i}^n.

Given the notion of an observation radius, we define *k-radius sets* (and similarly *perfect k-radius sets*) as follows.

Definition 4. *Let $V_k^{n,\sigma}$ be defined as*

$$V_k^{n,\sigma} = \{i \in \mathcal{N} \mid |B_{i,\tau_i^{n,\sigma}}^{n,\sigma}| \leq k\}.$$

We refer to $V_k^{n,\sigma}$ as the k-radius set.

Intuitively, $V_k^{n,\sigma}$ includes all agents that take an action before they receive signals from more than k other individuals at equilibrium σ^n - the size of their (indirect) neighborhood by the time they take an irreversible action is no greater than k. Equivalently, agent i belongs to set $V_k^{n,\sigma}$ if the number of agents that lie at distance less than $\tau_i^{n,\sigma}$ from i are at most k. From Definition 4 it follows immediately that

$$i \in V_k^{n,\sigma} \Rightarrow i \in V_{k'}^{n,\sigma} \text{ for all } k' > k. \tag{6.1}$$

The following proposition provides a necessary and a sufficient condition for ϵ, δ-asymptotic learning to occur in a society under equilibrium profile σ.

Proposition 15. *Let Assumption 6 hold. Then,*

(a) ϵ, δ-asymptotic learning occurs in society $\{G^n\}_{n=1}^{\infty}$ under equilibrium profile σ if

$$\lim_{n\to\infty} \frac{1}{n} \cdot |V_k^{n,\sigma}| < \zeta < \epsilon, \tag{6.2}$$

where $k > \hat{k}$ and $\hat{k}$ such that $erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/\hat{k}}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/\hat{k}}}\right) > 1 - \frac{\delta(\epsilon-\zeta)}{1-\zeta}$.

(b) ϵ, δ-asymptotic learning does not occur in society $\{G^n\}_{n=1}^{\infty}$ under equilibrium profile σ if

$$\lim_{n \to \infty} \frac{1}{n} \cdot \left| V_k^{n,\sigma} \right| > \eta > \epsilon, \tag{6.3}$$

where $k < \bar{k}$ and $\bar{k}$ such that $erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/\bar{k}}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/\bar{k}}}\right) < (1-\delta)(1-\epsilon/\eta)$.

Intuitively, asymptotic learning is precluded if there exists a significant fraction of the society that will take an action before seeing a large set of signals, since in this case there will be a large enough probability of each individual taking an action, that is far from the optimal one, since it is based on a small set of signals. In the rest of this section, we provide a series of corollaries of Proposition 15, that provide more intuition on the asymptotic learning result. The first one provides a *necessary* and sufficient condition for *perfect* asymptotic learning to occur in any equilibrium profile.

Before stating the proposition, we define the notion of *leading agents.* Let $indeg_i^n = |B_{i,1}^n|$, $outdeg_i^n = \left|\{j \middle| i \in B_{j,1}^n\}\right|$ denote the in-degree, out-degree of agent i in communication network G^n respectively.

Definition 5. *A collection $\{S\}_{n=1}^{\infty}$ of sets of agents is called a set of leading agents (leaders) if*

(i) There exists $k > 0$, such that $S^{n_j} \subseteq V_k^{n_j}$ for all $j \in J$, where J is an infinite index set.

(ii) $\lim_{n\to\infty} \frac{1}{n} \cdot \left|S^n\right| = 0$, i.e., the collection $\{S\}_{n=1}^{\infty}$ contains a negligible fraction of the agents as the society grows.

(iii) $\lim_{n\to\infty} \frac{1}{n} \cdot \left|S_{follow}^n\right| > \epsilon$, for some $\epsilon > 0$, where S_{follow}^n denotes the set of followers of S^n. In particular, $i \in S_{follow}^n$ if there exists $j \in S^n$, such that $j \in B_{i,\tau_i^n}^n$.

Then,

Proposition 16. *Let Assumption 6 hold. Then,*

(i) Perfect asymptotic learning occurs in society $\{G^n\}_{n=1}^{\infty}$ in any equilibrium σ if

$$\lim_{k\to\infty}\lim_{n\to\infty}\frac{1}{n}\cdot\left|V_k^n\right| = 0. \tag{6.4}$$

(ii) Conversely, if condition (6.4) does not hold for society $\{G^n\}_{n=1}^{\infty}$ and the society does not contain a set of leaders, then perfect asymptotic learning does not occur in any equilibrium σ.

Proposition 16 is not stated as an if and only if result because the fact that condition (6.4) does not hold in a society, does not necessarily preclude perfect asymptotic learning in the presence of leading agents. In particular, if the leaders delay their irreversible decision long enough, then a large fraction of the rest of the agents may take an irreversible action as soon as they communicate with the leading agents and, thus, perfect asymptotic learning fails (however the aggregate welfare is higher). However, if the leading agents do not coordinate at equilibrium, then they exit early and this may lead the rest of the agents to take a delayed, but more informed, irreversible action. Note that Proposition 16 is equilibrium independent, i.e., it holds for all equilibria and Condition (6.4) is expressed in terms of the network topology.

The next three corollaries identify the role of certain types of agents on information spread in a given society. We focus on perfect asymptotic learning, since we can obtain sharper results, however, we can state similar corollaries for ϵ, δ-asymptotic learning for any ϵ and δ. All corollaries are again expressed in terms of the original network topology.

Similarly, to perfect k-radius sets, we define sets U_k^n for scalar $k > 0$ as

$$U_k^n = \{i \in \mathcal{N}^n \,|\, \text{there exists } \ell \in B_{i,\tau_i^n}^n \text{ with } indeg_\ell^n > k\},$$

i.e., the set U_k^n consists of all agents, which have an agent with in-degree at least k within their perfect observation radius.

Corollary 2. *Let Assumption 6 hold. Then, perfect asymptotic learning occurs in society* $\{G^n\}_{n=1}^{\infty}$ *if*

$$\lim_{k\to\infty}\lim_{n\to\infty}\frac{1}{n}\cdot\left|U_k^n\right| = 1.$$

Intuitively, Corollary 2 states that if all but a negligible fraction of the agents are at a short distance (at most their perfect observation radius) from an agent with a high in-degree, then asymptotic learning occurs. This corollary therefore identifies a group of agents, that is crucial for a society to permit asymptotic learning: *information mavens*, who have high in-degrees and can thus act as effective aggregators of information (a term inspired by [36]). Information mavens are one type of hubs, the importance of which is clearly illustrated by our learning results. Our next definition formalizes this notion further and enables an alternative sufficient condition for asymptotic learning.

Definition 6. *Agent i is called an information maven of society $\{G^n\}_{n=1}^{\infty}$ if i has an infinite in-degree, i.e., if*

$$\lim_{n\to\infty} indeg_i^n = \infty.$$

Let $\mathcal{MAVEN}(\{G^n\}_{n=1}^{\infty})$ denote the set of mavens of society $\{G^n\}_{n=1}^{\infty}$.

For any agent j, let $d_j^{\mathcal{MAVEN},n}$ denote the shortest distance defined in communication network G^n between j and a maven $k \in \mathcal{MAVEN}(\{G^n\}_{n=1}^{\infty})$. Finally, let W^n denote the set of agents that are at distance at most equal to their perfect observation radius from a maven in communication network G^n, i.e., $W^n = \{j \mid\ d_j^{\mathcal{MAVEN},n} \leq \tau_j^n\}$.

The following corollary highlights the importance of information mavens for asymptotic learning. Informally, it states that if almost all agents have a short path to a maven, then asymptotic learning occurs.

Corollary 3. *Let Assumption 6 hold. Then, asymptotic learning occurs in society $\{G^n\}_{n=1}^{\infty}$ if*

$$\lim_{n\to\infty}\frac{1}{n}\cdot\left|W^n\right| = 1.$$

Corollary 3 thus clarifies that asymptotic learning is obtained when there are information mavens and almost all agents are at a "short distance" away from one (less than their observation radius).[1]

As mentioned in the Introduction, a second type of information hubs also plays an important role in asymptotic learning. While mavens have high in-degree and are thus able to effectively aggregate dispersed information, because our communication network is directed, they may not be in the right position to distribute this aggregated information. If so, even in a society that has several information mavens, a large fraction of the agents may not benefit from their information. *Social connectors*, on the other hand, are defined as agents with a high out-degree, and thus play the role of spreading the information aggregated by the mavens.[2] Before stating the proposition, we define *social connectors.*

Definition 7. *Agent i is called a social connector of society $\{G^n\}_{n=1}^{\infty}$ if i has an infinite out-degree, i.e., if*

$$\lim_{n\to\infty} outdeg_i^n = \infty.$$

The following corollary illustrates the role of social connectors for asymptotic learning.

Corollary 4. *Let Assumption 6 hold. Consider society $\{G^n\}_{n=1}^{\infty}$, which is such that the sequence of in- and out- degrees is non-decreasing for every agent and the set of information mavens does not grow at the same rate as the society itself, i.e.,*

$$\lim_{n\to\infty} \frac{|\mathcal{MAVEN}(\{G^n\}_{n=1}^{\infty})|}{n} = 0.$$

Then, for asymptotic learning to occur, the society should contain a social connector within

[1] This corollary is weaker than Corollary 2. This is simply a technicality because the sequence of communication networks $\{G^n\}_{n=1}^{\infty}$ is arbitrary. In particular, we have not assumed that the in-degree of an agent is non-decreasing with n, thus the limit in the corollary may not be well defined for arbitrary sequences of communication networks

[2] For simplicity (and to avoid the technical issues mentioned in a previous footnote) we assume for this corollary that both the in- and out-degree sequences of agents are non-decreasing with n (note that we can rewrite the proposition for any sequence of in- and out- degrees at the expense of introducing additional notation).

a short distance to a maven, i.e.,

$$d_i^{\mathcal{MAVEN},n} \leq \tau_i^n, \text{ for some social connector } i.$$

Corollary 4 thus states that unless a non-negligible fraction of the agents belongs to the set of mavens and, subsequently, the rest can obtain information directly from a maven, then, information aggregated at the mavens is spread through the out-links of a connector (note that an agent can be both a maven and a connector). Combined with the previous corollaries, this result implies that there are essentially two ways in which society can achieve perfect asymptotic learning. First, it may contain several information mavens who not only collect and aggregate information but also distribute it to almost all the agents in the society. Second, it may contain a sufficient number of information mavens, who pass their information to social connectors, and almost all the agents in the society are a short distance away from social connectors and thus obtain accurate information from them. This latter pattern has a greater plausibility in practice than one in which the same agents will collect and distribute dispersed information. For example, if a website or a news source can rely on information mavens (journalists, researchers or analysts) to collect sufficient information and then reach a large number of individuals, then information will be economically aggregated and distributed in the society.

The results summarized in Propositions 15 and 16 as well as in Corollaries 2, 3 and 4 can be seen both as positive and negative, as already noted in the Introduction. On the one hand, communication structures that do not feature information mavens (or connectors) will not lead to perfect asymptotic learning, and information mavens may be viewed as unrealistic or extreme. On the other hand, as already noted above, much communication in modern societies happens through agents that play the role of mavens and connectors (see again [36]). These are highly connected agents that are able to collect and distribute crucial

information. Perhaps more importantly, most individuals obtain some of their information from news sources, media, and websites, which exist partly or primarily for the purpose of acting as information mavens and connectors.[3]

6.3 Strategic Communication

Next we explore the implication of relaxing the assumption that agents cannot manipulate the messages they send, i.e., that information on private signals is hard. In particular, we replace Assumption 6 with Assumption 7 and we allow agents to lie about their or any of the private signals they have obtained information about, i.e., $m^n_{ij,t}(s_{i_1}, \cdots, s_{i_{|I^n_{i,t}|}}) \neq (s_{i_1}, \cdots, s_{i_{|I^n_{i,t}|}})$, where the latter is the true vector of private signals of agents in $I^n_{i,t}$. Informally, an agent has an incentive to misreport information, so as to delay her neighbors taking irreversible actions, which in turn prolongs the information exchange process.

Let (σ^n, m^n) denote an action-message strategy profile, where $m^n = \{m^n_1, \cdots, m^n_n\}$ and $m^n_i = [m^n_{ij,t}]_{t=0,1,\cdots}$, for j such that $i \in B^n_{j,1}$. Also let subscript (σ^n, m^n) refer to the probability measure induced by the action-message strategy profile.

Definition 8. *An action-message strategy profile $(\sigma^{n,*}, m^{n,*})$ is a pure-strategy Perfect Bayesian Equilibrium of the information exchange game $\Gamma_{info}(G^n)$ if for every $i \in \mathcal{N}^n$ and time t, we have*

$$\mathbb{E}_{(\sigma^{n,*}, m^{n,*})}(\pi^n_i \,|\, I^n_{i,t}, \sigma^{n,*}_{i,t}(I^n_{i,t})) \geq \mathbb{E}_{((\sigma^n_i, \sigma^{n,*}_{-i}),(m^n_i, m^{n,*}_{-i}))}(\pi^n_i \,|\, I^n_{i,t}, \sigma^n_{i,t}(I^n_{i,t})),$$

for all m^n_i and σ^n_i.

As before, we denote the set of equilibria of this game by $INFO(G^n, S^n)$.

Similarly we extend the definitions of asymptotic learning [cf. Definitions 2 and 3]. We show that strategic communication does not harm perfect asymptotic learning. The main

[3] For example, a news website such as cnn.com acts as a connector that spreads the information aggregated by the journalists-mavens to interested readers. Similarly, a movie review website, e.g., imdb.com, spreads the aggregate knowledge of movie reviewers to interested movie aficionados.

intuition behind this result is that it is weakly dominant for an agent to report her private signal truthfully to a neighbor with a high in-degree (maven), as long as others are truthful to the maven.

Theorem 1. *If perfect asymptotic learning occurs in society* $\{G^n\}_{n=1}^{\infty}$ *under Assumption 6, then there exists a equilibrium* (σ, m)*, such that perfect asymptotic learning occurs in society* $\{G^n\}_{n=1}^{\infty}$ *along equilibrium* (σ, m) *when we allow strategic communication (cf. under Assumption 7).*

This theorem therefore implies that the focus on truthful reporting was without much loss of generality as far as perfect asymptotic learning is concerned. In any communication network in which there is perfect asymptotic learning, even if agents can strategically manipulate information, there is arbitrarily little benefit in doing so. Thus, the main lessons about asymptotic learning derived above apply regardless of whether communication is strategic or not.

However, this theorem does not imply that all learning outcomes are identical under truthful and strategic communication. In particular, interestingly, we can construct examples in which strategic communication leads agents take a better action with higher probability than under non-strategic communication (cf. Assumption 6). The main reason for this (counterintuitive) fact is that under strategic communication an agent may delay taking an action compared to the non-strategic environment. Therefore, the agent obtains more information from the communication network and, consequently, chooses an action, that is closer to optimal.

6.4 Welfare

In this section, we turn attention to the question of *efficient* communication and draw a connection with our previous results on asymptotic learning. In particular, we compare equilibrium allocations, i.e., communication and action profiles at equilibrium, with a setting in which a social planner dictates the timing of agents' actions. We identify conditions

under which a social planner can / cannot improve in terms of the aggregate expected welfare over an equilibrium strategy profile. In doing so, we illustrate an interesting feature of social networks: communication might be inefficient, because agents do not internalize the positive externality that delaying their action generates for their peers.

A social planner, whose objective is to maximize the aggregate expected welfare of the population of n agents, would choose to implement the timing profile, that is obtained as a solution to the optimization problem below. We call this timing profile the *optimal allocation* and we denote it by sp^n. Note that $sp^n = (\tau_1^n, \cdots, \tau_n^n)$

$$\max_{sp^n} \sum_{i=1}^{n} \mathbb{E}_{sp^n}[u_i^n] \tag{6.5}$$

Similarly with the asymptotic analysis for equilibria, we define a sequence of optimal allocations for societies of growing size, $sp = \{sp^n\}_{n=1}^{\infty}$. We are interested in identifying conditions under which the social planner can / cannot achieve an *asymptotically* better allocation than an equilibrium (sequence of equilibria) σ, i.e., we are looking at the expression:

$$\lim_{n\to\infty} \frac{\sum_{i\in N^n} \mathbb{E}_{sp^n}[u_i^n] - \sum_{i\in N^n} \mathbb{E}_{\sigma}[u_i]}{n}.$$

First, we state a theorem that shows a direct connection between learning and efficient communication.

Theorem 2. *Consider society $\{G^n\}_{n=1}^{\infty}$. If perfect asymptotic learning occurs at the optimal allocation $sp = \{sp^n\}_{n=1}^{\infty}$, then all equilibria are asymptotically efficient, i.e.,*

$$\lim_{n\to\infty} \frac{\sum_{i\in N^n} \mathbb{E}_{sp^n}[u_i^n] - \sum_{i\in N^n} \mathbb{E}_{\sigma}[u_i]}{n} = 0,$$

for all equilibria σ.

Note that from Proposition 16 it follows that if perfect learning occurs at the optimal

allocation, then perfect learning occurs in all equilibria σ.

In the remainder of the section, we provide a result that can be thought of as a partial converse of Theorem 2. Before stating it, we contrast the single agent decision problem with that of the social planner. In particular, consider agent i. Then, i will decide to take an irreversible action at time t and not wait for an additional dt, when other agents behave according to σ, if (cf. Section 6.1):

$$\frac{r+\lambda}{r}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k_{i,t}^{n,\sigma}}\right) \geq V_i(k_{i,t}^{n,\sigma} + |B_{i,|T_t|+1}^{n,\sigma}| - |B_{i,|T_t|}^{n,\sigma}|, \sigma) \tag{6.6}$$

Similarly, in the corresponding optimal allocation agent i will exit at time t and not wait if:

$$\frac{r+\lambda}{r}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k_{i,t}^{n,sp}}\right)$$
$$\geq V_i(k_{i,t}^{n,sp} + |B_{i,|T_t|+1}^{n,sp}| - |B_{i,|T_t|}^{n,sp}|, sp) + \sum_{j\neq i} \mathbb{E}_{sp}[u_j^n | i \text{ "waits" at } t] - \mathbb{E}_{sp}[u_j^n | i \text{ "exits" at } t]. \tag{6.7}$$

Contrasting (6.6) with (6.7) shows the reason for why equilibria are inefficient in this setting: when determining when to act agent i does not take into account the positive externality that a later action exerts on others, which is expressed by the summation on the right hand side of (6.7). We are left with showing sufficient conditions under which a social planner outperforms an equilibrium allocation σ. Consider agents i and j such that $i \in B_{j,1}^n$ and $\tau_j^{n,\sigma} > \tau_i^{n,\sigma} + 1$, which implies that $B_{j,\tau_j^{n,\sigma}}^n \supset B_{i,\tau_i^{n,\sigma}}^n$ (i.e., agent j communicates with a superset of the agents that i communicates with before taking an action). Also, let $k_{ij,\tau_i^{n,\sigma}}^{n,\sigma}$ denote the additional agents that j would observe if i delayed her irreversible action by dt and communication took place. Then, the aggregate welfare of the two agents increases if

the following condition holds:

$$V_j(k^{n,\sigma}_{j,\tau^{n,\sigma}_j} + k^{n,\sigma}_{ij,\tau^{n,\sigma}_i}) + V_i(k^{n,\sigma}_{i,\tau^{n,\sigma}_i} + k^{n,\sigma}_{ij,\tau^{n,\sigma}_i}) > V_j(k^{n,\sigma}_{j,\tau^{n,\sigma}_j}) + \frac{r+\lambda}{\lambda} V_i(k^{n,\sigma}_{i,\tau^{n,\sigma}_i}), \tag{6.8}$$

In particular, let $k^{n,\sigma}_{j,\tau^{n,\sigma}_j} \leq \bar{k}$ (and in consequence $k^{n,\sigma}_{i,\tau^{n,\sigma}_i} \leq \bar{k}$). Then, (6.8) holds if $k_1 < k^{n,\sigma}_{ij,\tau^{n,\sigma}_j} < k_2$ (k_1, k_2 constants), where the lower bound guarantees that the aggregate welfare increases by the additional delay, whereas the upper bound ensures that agent i would find it optimal to "exit" after communication step $\tau^{n,\sigma}_i$ at equilibrium σ.

Now let set $D^{n,\sigma}_{k,\ell}$ denote the following set of agents: $j \in D^{n,\sigma}_{k,\ell}$, if $k^{n,\sigma}_{j,\tau^{n,\sigma}_j} \leq k$ and there exists an $i \in B^{n,\sigma}_{j,\tau^{n,\sigma}_j}$ such that if i exits at $\tau^{n,\sigma}_i + 1$, then j gains access to at least an additional ℓ signals. Then, the following proposition provides a sufficient condition for an equilibrium to be inefficient.

Proposition 17. *Consider society $\{G^n\}_{n=1}^{\infty}$ and equilibrium $\sigma = \{\sigma^n\}_{n=1}^{\infty}$. Assume that*

$$\lim_{n\to\infty} \frac{|D^{n,\sigma}_{k,\ell}|}{n} > \xi > 0,$$

for k, ℓ that satisfy the following:

$$\frac{2}{1/\rho^2 + 1/\bar{\rho}^2(k+\lambda)} < \left(2 + \frac{r}{\ell}\right) \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k} + \frac{r}{\lambda}\pi$$

Then, there exists an $\zeta > 0$, such that

$$\lim_{n\to\infty} \frac{\sum_{i\in N^n} \mathbb{E}_{sp^n}[u^n_i] - \sum_{i\in N^n} \mathbb{E}_{\sigma}[u_i]}{n} > \zeta,$$

i.e., equilibrium σ is asymptotically inefficient. Moreover, there exist ϵ, δ such that ϵ, δ-asymptotic learning fails at equilibrium σ.

Chapter 7

Asymptotic Learning in Random Graphs

As an illustration of the results we outlined in Section 6, we apply them to a series of commonly studied random graph models. We begin by providing the definitions for the graph models we focus on. Note that in the present section we assume that communication networks are bidirectional, or equivalently that if agent $i \in B_{j,1}^n$ then $j \in B_{i,1}^n$.

Definition 9. *A sequence of communication networks $\{G^n\}_{n=1}^{\infty}$, where $G^n = \{\mathcal{N}^n, \mathcal{E}^n\}$, is called*

(i) complete if for every n we have

$$(i,j) \in \mathcal{E}^n \quad \textit{for all } i,j \in \mathcal{N}^n \ .$$

(ii) k-bounded degree for scalar $k > 0$, if for every n we have

$$|B_{i,1}^n| \leq k \quad \textit{for all } i \in \mathcal{N}^n \ ,$$

where recall that $B_{i,1}^n$ denotes the agents that are one link away from agent i in com-

munication network G^n.

(iii) a star if for every n we have

$$(i,1) \in \mathcal{E}^n \text{ and } (i,j) \notin \mathcal{E}^n \quad \text{for all } i \in \mathcal{N}^n \text{ and } j \neq 1 \, .$$

Definition 10 (Erdős-Renyi). *A sequence of communication networks $\{G^n\}_{n=1}^{\infty}$, where $G^n = \{\mathcal{N}^n, \mathcal{E}^n\}$, is called Erdős-Renyi if for every n we have*

$$\mathbb{P}\left((i,j) \in \mathcal{E}^n\right) = \frac{p}{n} \quad \text{independently for all } i,j \in \mathcal{N}^n \ ,$$

where p scalar, such that $0 < p < 1$.

Definition 11 (Power-Law). *A sequence of communication networks $\{G^n\}_{n=1}^{\infty}$, where $G^n = \{\mathcal{N}^n, \mathcal{E}^n\}$, is called Power-Law with exponent $\gamma > 0$ if we have*

$$\lim_{n\to\infty} \frac{\sum_{i\in\mathcal{N}^n} \mathbf{1}_{|B_{i,1}^n|=k}}{n} = c_k \cdot k^{-\gamma} \text{ for every scalar } k > 0,$$

where c_k is a constant. In other words, the fraction of nodes in the network having degree k, for every $k > 0$, follows a power law distribution with exponent γ.

Definition 12 (Preferential Attachment). *A sequence of communication networks $\{G^n\}_{n=1}^{\infty}$, where $G^n = \{\mathcal{N}^n, \mathcal{E}^n\}$, is called preferential attachment if it was generated by the following process:*

(i) Begin the process with $G^1 = \{\{1\}, \{(1,1)\}\}$, i.e., the communication network that contains agent 1 and a loop edge.

(ii) At step n, add agent $n+1$ to G^n. Choose an agent w from G^n and let $\mathcal{E}^{n+1} = \mathcal{E}^n + (n+1, w)$. Agent w is chosen according to the preferential attachment rule, i.e.,

$w = j$ *for* $j \in N^n$ *with probability*

$$\mathbb{P}(w = j) = \frac{deg(j)}{\sum_{\ell \in \mathcal{N}^n} deg(\ell)},$$

where $deg(j)$ *denotes the degree of node* j *at the step.*

Definition 13 (Hierarchical). *A sequence of communication networks* $\{G^n\}_{n=1}^{\infty}$, *where* $G^n = \{\mathcal{N}^n, \mathcal{E}^n\}$, *is called* ζ*-hierarchical (or simply hierarchical) if it was generated by the following process:*

(i) Agents are born and placed into layers. In particular, at each step $n = 1, \cdots$, *a new agent is born and placed in layer* ℓ.

(ii) Layer index ℓ *is initialized to 1 (i.e., the first node belongs to layer 1). A new layer is created (and subsequently the layer index increases by one) at time period* $n \geq 2$ *with probability* $\frac{1}{n^{1+\zeta}}$, *where* $\zeta > 0$.

(iii) Finally, for every n *we have*

$$\mathbb{P}\left((i,j) \in \mathcal{E}^n\right) = \tfrac{p}{|\mathcal{N}_\ell^n|}, \text{ independently for all } i, j \in \mathcal{N}^n \text{ that belong to the same layer } \ell,$$

where $\mathcal{N}_\ell^n$ *denotes the set of agents that belong to layer* ℓ *at step* n *and* p *scalar, such that* $0 < p < 1$. *Moreover,*

$$\mathbb{P}\left((i,k) \in \mathcal{E}^n\right) = \frac{1}{|\mathcal{N}_{<\ell}|} \text{ and } \sum_{k \in \mathcal{N}_{<\ell}} \mathbb{P}\left((i,k) \in \mathcal{E}^n\right) = 1 \text{ for all } i \in N_\ell^n, k \in N_{<\ell}^n, \ell > 1,$$

where $\mathcal{N}_{<\ell}^n$ *denotes the set of agents that belong to a layer with index lower than* ℓ *at step* n.

Intuitively, a hierarchical sequence of communication networks resembles a pyramid, where the top contains only a few agents and as we move towards the base, the number of agents grows. The following argument provides an interpretation of the model. Agents

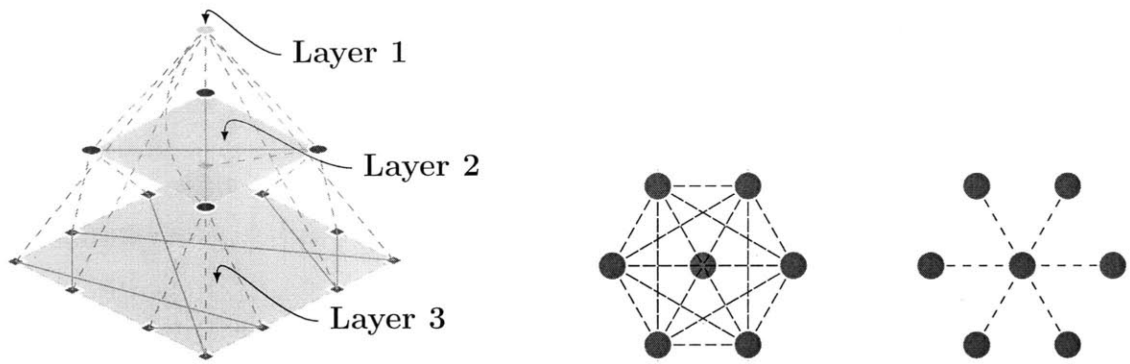

(a) Hierarchical Society. (b) Complete Society. (c) Star Society.

Figure 7-1: Example society structures.

on top layers can be thought of as "special" nodes, that the rest of the nodes have a high incentive connecting to. Moreover, agents tend to connect to other agents in the same layer, as they share common features with them (*homophily*). As a concrete example, academia can be thought of as such a pyramid, where the top layer includes the few institutions, then next layer includes academic departments, research labs and finally at the lower levels reside the home pages of professors and students.

Proposition 18. *Let Assumptions 6 hold and consider society* $\{G^n\}_{n=1}^{\infty}$ *and discount rate* $r > 0$. *Then,*

(i) Perfect asymptotic learning does not occur in society $\{G^n\}_{n=1}^{\infty}$ *if the sequence of communication networks* $\{G^n\}_{n=1}^{\infty}$ *is* k*-bounded, for some constant* $k > 0$.

(ii) For every $\epsilon > 0$, *asymptotic learning does not occur in society* $\{G^n\}_{n=1}^{\infty}$ *with probability at least* $1 - \epsilon$, *if the sequence of communication networks* $\{G^n\}_{n=1}^{\infty}$ *is preferential attachment.*

(iii) For every $\epsilon > 0$, *asymptotic learning does not occur in society* $\{G^n\}_{n=1}^{\infty}$ *with probability at least* $1 - \epsilon$, *if the sequence of communication networks* $\{G^n\}_{n=1}^{\infty}$ *is Erdős-Renyi.*

Proposition 19. *Let assumptions 6 hold and consider society* $\{G^n\}_{n=1}^{\infty}$. *Then,*

(i) Asymptotic learning occurs in society $\{G^n\}_{n=1}^{\infty}$ *if the sequence of communication networks* $\{G^n\}_{n=1}^{\infty}$ *is complete and the discount rate* r *is smaller than some scalar* $r_1 > 0$.

(ii) *Asymptotic learning occurs in society* $\{G^n\}_{n=1}^{\infty}$ *if the sequence of communication networks* $\{G^n\}_{n=1}^{\infty}$ *is a star and the discount rate* r *is smaller than some scalar* $r_2 > 0$.

(iii) *Let* $\epsilon > 0$. *Then, asymptotic learning occurs in society* $\{G^n\}_{n=1}^{\infty}$ *with probability at least* $1 - \epsilon$, *if the sequence of communication networks* $\{G^n\}_{n=1}^{\infty}$ *is* γ*-power law, with* $\gamma \leq 2$ *and the discount rate* r *is smaller than some scalar* $r_3 > 0$.

(iv) *Let* $\epsilon > 0$. *Then, asymptotic learning occurs in society* $\{G^n\}_{n=1}^{\infty}$ *with probability at least* $1 - \epsilon$, *if the sequence of communication networks* $\{G^n\}_{n=1}^{\infty}$ *is* $\zeta(\epsilon)-$*hierarchical and the discount rate* r *is smaller than some scalar* $r_4 > 0$.

The results presented provide additional insights on the conditions under which asymptotic learning takes place. The popular preferential attachment and Erdős-Renyi graphs do not lead to asymptotic learning, which can be interpreted as implying that asymptotic learning is unlikely in several important networks. Nevertheless, these network structures, though often used in practice, do not provide a good description of the structure of many real life networks. In contrast, our results also showed that asymptotic learning takes place in power law graphs with small exponent $\gamma \leq 2$, and such graphs appear to provide a better representation for many networks related to communication, including for peer-to-peer networks. scientific collaboration networks (in experimental physics) and traffic in networks ([50], [57], [69], [70]). Asymptotic learning also takes place in hierarchical graphs, where "special" agents are likely to receive and distribute information to lower layers of the hierarchy.

Chapter 8

Network Formation

In previous chapters, we studied information exchange among agents over a given communication network $G^n = (\mathcal{N}^n, \mathcal{E}^n)$. In the present chapter, we turn attention to the complementary problem of how this communication network emerges. We assume that link formation is costly. In particular, *communication costs* are captured by an $n \times n$ non-negative matrix C^n, where C^n_{ij} denotes the cost that agent i has to incur in order to form the directed link (j, i) with agent j. As noted previously, a link's direction coincides with the direction of the flow of messages. In particular, agent i incurs a cost to form in-links. We refer to C^n as the *communication cost matrix.* We assume that $C^n_{ii} = 0$ for all $i \in \mathcal{N}^n$.

We define agent i's *link formation strategy*, g^n_i, as an n-tuple such that $g^n_i \in \{0, 1\}^n$ and $g^n_{ij} = 1$ implies that agent i forms a link with agent j. The cost agent i has to incur if she implements strategy g^n_i is given by

$$\text{Cost}(g^n_i) = \sum_{j \in \mathcal{N}} C^n_{ij} \cdot g^n_{ij}.$$

The link formation strategy profile $g^n = (g^n_1, \cdots, g^n_n)$ induces the communication network $G^n = (\mathcal{N}^n, \mathcal{E}^n)$, where $(j, i) \in \mathcal{E}^n$ if and only if $g^n_{ij} = 1$.

We extend the previously described environment to the two-stage *Network Learning*

Game $\Gamma(C^n)$, where C^n denotes the communication cost matrix. The two stages of the network learning game can be described as follows:

Stage 1 [*Network Formation Game*]: Agents pick their link formation strategies. The link formation strategy profile g^n induces the communication network $G^n = (\mathcal{N}^n, \mathcal{E}^n)$. We refer to stage 1 of the network learning game, when the communication cost matrix is C^n as the *network formation game* and we denote it by $\Gamma_{net}(C^n)$.

Stage 2 [*Information Exchange Game*]: Agents communicate over the induced network G^n as studied in previous chapters.

We next define the equilibria of the network learning game $\Gamma(C^n)$. Note that we use the standard notation g_{-i} and σ_{-i} to denote the strategies of agents other than i. Also, we let $\sigma_{i,-t}$ denote the vector of actions of agent i at all times except t.

Definition 14. *A pair $(g^{n,*}, \sigma^{n,*})$ is a pure-strategy Perfect Bayesian Equilibrium of the network learning game $\Gamma(C^n)$ if*

(a) $\sigma^{n,*} \in INFO(G^n)$, where G^n is induced by the link formation strategy $g^{n,*}$.

(b) For all $i \in \mathcal{N}^n$, $g_i^{n,*}$ maximizes the expected payoff of agent i given the strategies of other agents $g_{-i}^{n,*}$, i.e.,

$$g_i^{n,*} \in \arg\max_{z \in \{0,1\}^n} \mathbb{E}_\sigma[\Pi_i(z, g_{-i}^{n,*})] \equiv \mathbb{E}_\sigma(u_i^n \big| I_{i,0}^n) - Cost(z).$$

for all $\sigma \in INFO(\tilde{G}^n)$, where $\tilde{G}^n$ is induced by link formation strategy $(z, g_{-i}^{n,*})$.

We denote the set of equilibria of this game by $NET(C^n)$.

8.1 Learning in endogenous networks

Similar to the analysis of the information exchange game, we consider a sequence of communication cost matrices $\{C^n\}_{n=1}^{\infty}$, where for fixed n,

$$C^n : \mathcal{N}^n \times \mathcal{N}^n \to \Re^+ \text{ and } C_{ij}^n = C_{ij}^{n+1} \text{ for all } i, j \in \mathcal{N}^n. \tag{8.1}$$

For the remainder of the section, we focus our attention to the *social cliques communication cost structure.* The properties of this communication structure are stated in the next assumption.

Assumption 8. *Let $c_{ij}^n \in \{0, c\}$ for all pairs $(i,j) \in \mathcal{N}^n \times \mathcal{N}^n$, where $c < \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}$. Moreover, let $c_{ij} = c_{ji}$ for all $i, j \in \mathcal{N}^n$ (symmetry), and $c_{ij} + c_{jk} \geq c_{ik}$ for all $i, j, k \in \mathcal{N}^n$ (triangular inequality).*

The assumption that $c < \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}$ rules out the degenerate case where no agent forms a costly link. The symmetry and triangular inequality assumptions are imposed to simplify the definition of a social clique, which is introduced next. Let Assumption 8 hold. We define a *social clique* (cf. Figure 8-1) $H^n \subset \mathcal{N}^n$ as a set of agents such that

$$i, j \in H^n \quad \text{if and only if} \quad c_{ij} = c_{ji} = 0.$$

Note that this set is well-defined since, by the triangular inequality and symmetry assumptions, if an agent i does not belong to social clique H^n, then $c_{ij} = c$ for all $j \in H^n$. Hence, we can uniquely partition the set of nodes $\mathcal{N}^n$ into a set of K^n pairwise disjoint social cliques $\mathcal{H}^n = \{H_1^n, \cdots, H_{K^n}^n\}$. We use the notation $\mathcal{H}_k^n$ to denote the set of pairwise disjoint social cliques that have cardinality greater than or equal to k, i.e., $\mathcal{H}_k^n = \{H_i^n, i = 1, \ldots, K^n \mid |H_i^n| \geq k\}$. We also use $sc^n(i)$ to denote the social clique that agent i belongs to.

We consider a sequence of communication cost matrices $\{C^n\}_{n=1}^{\infty}$ satisfying condition (8.1) and Assumption 8, and we refer to this sequence as a *communication cost structure.* As shown above, the communication cost structure $\{C^n\}_{n=1}^{\infty}$ uniquely defines the following sequences, $\{\mathcal{H}^n\}_{n=1}^{\infty}$ and $\{\mathcal{H}_k^n\}_{n=1}^{\infty}$ for $k > 0$, of sets of pairwise disjoint social cliques. Moreover, it induces network equilibria $(g, \sigma) = (g^n, \sigma^n)_{n=1}^{\infty}$ such that $(g^n, \sigma^n) \in NET(C^n)$ for all n. Our goal is to identify conditions on the communication cost structure, that lead to the emergence of networks, that guarantee asymptotic learning. We focus entirely on

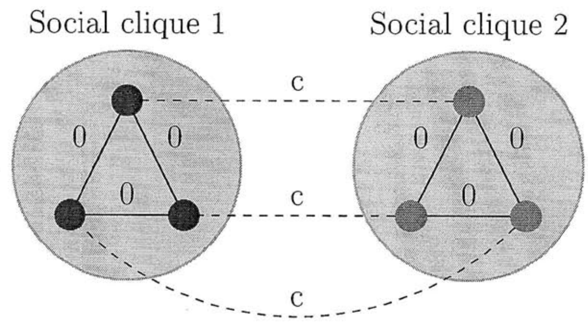

Figure 8-1: Social cliques.

perfect asymptotic learning, as this enables us to obtain sharp results.

Theorem 3. *Let $\{C^n\}_{n=1}^{\infty}$ be a communication cost structure and let Assumptions 6 and 8 hold. Then, there exists a constant $\bar{k} = \bar{k}(c)$ such that the following hold:*

(a) Suppose that

$$\limsup_{n\to\infty} \frac{|\mathcal{H}_{\bar{k}}^n|}{n} \geq \epsilon \text{ for some } \epsilon > 0. \tag{8.2}$$

Then, perfect asymptotic learning does not occur in any network equilibrium (g, σ).

(b) Suppose that

$$\lim_{n\to\infty} \frac{|\mathcal{H}_{\bar{k}}^n|}{n} = 0 \text{ and } \lim_{n\to\infty} |H_\ell^n| = \infty \text{ for some } \ell. \tag{8.3}$$

Then, perfect asymptotic learning occurs in all network equilibria (g, σ) when the discount rate r satisfies $0 < r < \bar{r}$, where $\bar{r} > 0$ is a constant.

(c) Suppose that

$$\lim_{n\to\infty} \frac{|\mathcal{H}_{\bar{k}}^n|}{n} = 0 \text{ and } \limsup_{n\to\infty} |H_\ell^n| < M \text{ for all } \ell, \tag{8.4}$$

where $M > 0$ is a scalar, and let agents be patient, i.e., consider the case, when the discount rate $r \to 0$. Then, there exists a $\bar{c} > 0$ such that

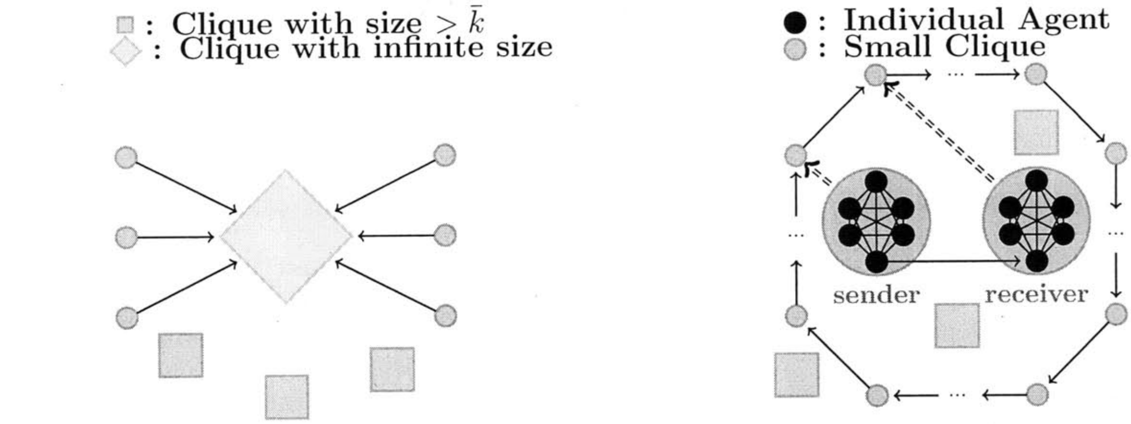

(a) Equilibrium network, when (8.3) holds. (b) Equilibrium network, when (8.4) holds.

Figure 8-2: Network formation among social cliques.

(i) If $c \leq \bar{c}$*, perfect asymptotic learning occurs in all network equilibria* (g, σ)*.*

(ii) If $c > \bar{c}$*, perfect asymptotic learning depends on the network equilibrium considered.*

In particular, there exists at least one network equilibrium (g, σ)*, where there is no perfect asymptotic learning and there exists at least one network equilibrium* (g, σ) *where perfect asymptotic learning occurs.*

The results in this theorem provide a fairly complete characterization of what types of environments will lead to the formation of networks that will subsequently induce perfect asymptotic learning. The key concept is that of a *social clique*, which represents groups of individuals that are linked to each other at zero cost. These can be thought of as "friendship networks," which are linked for reasons unrelated to information exchange and thus can act as conduits of such exchange at low cost. Agents can exchange information without incurring any costs (beyond the delay necessary for obtaining information) within their social cliques. However, if they wish to obtain further information, from outside their social cliques, they have to pay a cost at the beginning in order to form a link. Even though network formation games have several equilibria, the structure of our network formation and information exchange game enables us to obtain relatively sharp results on

what types of societies will lead to endogenously formed communication networks that ensure perfect asymptotic learning. In particular, the first part of Theorem 3 shows that perfect asymptotic learning cannot occur in any equilibrium if the number of sufficiently large social cliques increases at the same rate as the size of the society (or its lim sup does so). This is intuitive; when this is the case, there will be many social cliques of sufficiently large size that none of their members wish to engage in further costly communication with members of other social cliques. But since several of these will not contain an information hub social learning is precluded.

In contrast, the second part of the theorem shows that if the number of disjoint and sufficiently large social cliques is limited (grows less rapidly than the size of the society) and some of them are large enough to contain information hubs, then perfect asymptotic learning will take place (provided that future is not heavily discounted). In this case, as shown by Figure 8-2(a), sufficiently many social cliques will connect to the larger social cliques acting as information hubs, ensuring effective aggregation of information for the great majority of the agents in the society. It is important that the discount factor is not too small, otherwise smaller cliques will not find it beneficial to form links with the larger cliques.

Finally, the third part of the theorem outlines a more interesting configuration, potentially leading to perfect asymptotic learning. In this case, many small social cliques form an "informational ring"(Figure 8-2(b)) . Each is small enough that it finds it beneficial to connect to another social clique, provided that this other clique will also connect to others and obtain further information. This intuition also clarifies why such information aggregation takes place only in some equilibria. The expectation that others will not form the requisite links leads to a coordination failure. Interestingly, however, if agents are sufficiently patient and the cost of link formation is not too large, the coordination failure equilibrium disappears, because it becomes beneficial for each clique to form links with another one, even if further links are not forthcoming.

Chapter 9

Conclusion

This thesis studies optimal stopping problems, when agents interact with each other in a strategic fashion. We describe our framework in the context of two economic applications: (1) experimentation and innovation in a competitive multi-firm R&D market and (2) information exchange among individuals, that are embedded in a social network structure. The key feature of our framework is a tradeoff that economic agents are facing when choosing when to act. Observing other agents' actions (and their outcomes) may increase one's expected payoff either directly (experimentation model) or indirectly (through information spillovers). Thus, there is an incentive to delay. However, future is discounted, therrefore agents prefer earlier than later actions. We conclude by summarizing the results we obtain in both models. We also discuss a number of interesting avenues for future research.

9.1 Experimentation, patents and innovation

In the first part of the thesis, we present a model of experimentation and innovation in a multi-firm environment. In the baseline version of the model, each firm receives a private signal on the success probability of one of many potential research projects and decides when and which project to implement. A successful innovation can be copied by other firms, thus reducing the profits of the original innovator. Firms have an incentive to delay their experimentation decision (investment in R&D) and free-ride on the efforts

of their competitors. Naturally, symmetric equilibria, where actions do not depend on the identity of the firm, necessarily involve delayed and staggered experimentation. When the signal quality is the same for all players, the equilibrium is in mixed strategies (pure-strategy symmetric equilibria do not exist). When signal quality differs across firms, the equilibrium is represented by a function $\tau(p)$ which specifies the time at which a firm with signal quality p experiments. As in the environment with common signal quality, the equilibrium may involve arbitrarily long delays.

We also show that the social cost of insufficient experimentation incentives can be arbitrarily large. The optimal allocation may require simultaneous rather than staggered experimentation. In this case, the efficiency gap between the optimal allocation and the equilibrium can be arbitrarily large. Instead, when the optimal allocation also calls for staggered experimentation, the equilibrium is inefficient because of delays. We show that in this case the ratio of social surplus in the equilibrium to that in the optimal allocation can be as low as 1/2.

One of the main arguments we make is that appropriately-designed patent systems encourage experimentation and reduce delays without preventing efficient *ex post* transfer of knowledge across firms. Consequently, when signal quality is the same for all firms, an appropriately-designed patent system can ensure that the optimal allocation results in all equilibria. Patents are particularly well-suited to providing the correct incentives when the optimal allocation also requires staggered experimentation. In this case, patents can simultaneously encourage one of the firms to play the role of a leader in experimentation, while providing incentives to others to copy successful innovations. Technically, appropriately-designed patents destroy symmetric equilibria, which are the natural equilibria in the absence of patents but may involve a high degree of inefficiency. That patents are an attractive instrument in this environment can also be seen from our result that, while patents can implement the optimal allocation, there exists no simple subsidy (to experimentation, research, or innovation) that can achieve the same policy objective. When

signal quality differs across firms, patents are again useful in encouraging experimentation and reduce delays, however typically they are unable to ensure the optimal allocation.

We believe that the role of patents in encouraging socially beneficial experimentation is more general than the simple model used in this thesis. In particular, throughout our analysis we only briefly consider the consumer side. It is possible that new innovations create benefits to consumers that are disproportionately greater than the use of existing successful innovations (as compared to the relative profitabilities of the same activities). In this case, the social benefits of experimentation are even greater and patents can also be useful in preventing copying of previous successful innovations. The investigation of the welfare and policy consequences of pursuing successful lines versus experimenting with new, untried research lines is an interesting and underresearched area.

9.2 Information exchange in endogenous social networks

In the second part of the thesis, we develop a model of information exchange through communication in social networks and investigate its implications for information aggregation in large societies. An underlying state (of the world) determines which action has higher payoff. Agents decide which agents to form a communication link with incurring the associated cost and receive a private signal correlated with an underlying state. They then exchange information over their communication network until taking an (irreversible) action.

Our focus has been on asymptotic learning, defined as the fraction of agents taking a nearly optimal action converging to one in probability as a society grows large. We show that perfect asymptotic learning occurs if and, under some additional mild assumptions, only if the communication network includes information hubs and most agents are at a short distance from a hub. Thus asymptotic learning requires information to be aggregated in the hands of a few agents. This kind of aggregation also requires truthful communication,

which we show is an equilibrium of the strategic communication in large societies. We also explore the welfare implications of equilibrium behavior by comparing the aggregate expected welfare at equilibrium with that at the optimal allocation (when a social planner dictates the timing of actions).

In the third part, we extend the information exchange model by allowing endogenous network formation. Using our analysis of information exchange over a given network, we provide a systematic investigation of what types of cost structures, and associated social cliques which consist of groups of individuals linked to each other at zero cost (such as friendship networks), ensure the emergence of communication networks that lead to asymptotic learning. Our main result on network formation shows that societies with too many (disjoint) and sufficiently large social cliques do not form communication networks that lead to asymptotic learning, because each social clique would have sufficient information to make communication with others not sufficiently attractive. Asymptotic learning results if social cliques are neither too numerous nor too large so as to encourage communication across cliques.

Interesting avenues for research include investigation of similar dynamic models of information exchange and network formation in the presence of ex ante or ex post heterogeneity of preferences as well as differences in the quality of information available to different agents, which may naturally lead to the emergence of hubs. Furthermore, the present model considers only purely informational externalities among the agents. It is conceivable, though, that one's action directly affects the expected payoffs of her peers. Assuming that agents' incentives are misaligned leads to the following natural question: under what conditions can truthful communication be sustained? Alternatively, one can think of introducing payoff externalities in our model as extending cheap-talk models by considering a multi-period, dynamic environment (as opposed to the static, one-shot models, that are prevalent in the literature).

Appendix A

Omitted Proofs from Chapters 1-4

Proof of Lemma 1

The proof comprises three steps. First, we show that $t = 0$ belongs to the support of mixing time (so that there is no time interval with zero probability of experimentation). Suppose, to obtain a contradiction, that $t_1 = \inf\{t\colon \lambda(t) > 0\} > 0$. Then, because experimenting after t_1 is in the support of the mixed-strategy equilibrium, equilibrium payoffs must satisfy

$$V_1 = e^{-rt_1} p \Pi_2.$$

Now consider deviation where firm i chooses $\lambda(0) = +\infty$. This has payoff

$$V_0 = p\Pi_2 > V_1$$

for any $t_1 > 0$, yielding a contradiction.

Second, we show that there does not exist $T < \infty$ such that the support of the stopping time τ (induced by λ) is within $[0, T]$. Suppose not, then it implies that there exists $t \in [0, T]$ such that $\lambda(t) = +\infty$ and let $t_1 = \inf\{t\colon \lambda(t) = +\infty\}$. This implies that the payoff to both firms once the game reaches time t_1 without experimentation (which has

positive probability since $t_1 = \inf\{t\colon \lambda(t) = +\infty\}$) is

$$V(\tau = t_1) = e^{-rt_1} p\Pi_2$$

(where $V(\tau = t)$, or $V(t)$, denotes present discounted value as a function of experimentation time). Now consider a deviation by firm i to strategy τ', which involves waiting for $\epsilon > 0$ after the game has reached t_1 and copying a successful project by firm $\sim i$ (if there is such a success). This has payoff

$$V(\tau') = e^{-r(t+\epsilon)} \left[p\Pi_2 + (1-p)p\Pi_1\right]$$

since firm $\sim i$ is still $\lambda(t_1) = +\infty$ and will thus experiment with probability 1 at t_1. Assumption 1 implies that $V(\tau')$ is strictly greater than $V(\tau = t_1)$ for ϵ sufficiently small.

Finally, we show that $\lambda(t) > 0$ for all t. Again suppose, to obtain a contradiction, that there exist t_1 and $t_2 > t_1$ such that $\lambda(t) = 0$ for $t \in (t_1, t_2)$. Then, with the same argument as in the first part, the payoff from the candidate equilibrium strategy τ to firm i conditional on no experimentation until t_1 is

$$V(\tau) = e^{-rt_2} p\Pi_2.$$

However, deviating and choosing $\tau' = t_1$ yields

$$V(\tau' = t_1) = e^{-rt_1} p\Pi_2 > V(\tau).$$

This contradiction completes the proof of the lemma.

Proof of Proposition 12

The proof consists of two main steps. The first involves characterizing the equilibrium with two firms when p has support $[a, b] \subset [0, \beta]$. The second involves extending the

characterization of equilibrium to the more general case when the support of G is $[a,b] \subset [0,1]$.

Step 1: We show that under the assumption that $[a,b] \subset [0,\beta]$, there exists a symmetric equilibrium represented by a strictly decreasing function $\tau(p)$ with $\tau(b) = 0$ which maps signals to time of experimentation provided that the other player has not yet experimented. Proposition 20 formalizes this idea and is proved by using a series of lemmas.

Proposition 20. *Suppose that the support of G is $[a,b] \subset [0,\beta]$. Define*

$$\tau(p) = \frac{1}{r\beta}\left[\log G(b)(1-b) - \log G(p)(1-p) + \int_p^b \log G(z)dz\right]. \tag{A.1}$$

Then the unique symmetric equilibrium takes the following form:

1. *each firm copies a successful innovation and immediately experiments if the other firm experiments unsuccessfully;*

2. *firm i with signal quality p_i experiments at time $\tau(p_i)$ given by (A.1) unless firm $\sim i$ has experimented before time $\tau(p_i)$.*

Proof. The proof uses the following lemmas.

Lemma 5. *$\tau(p)$ cannot be locally constant. That is, there exists no interval $P = [\bar{p}, \bar{p}+\epsilon]$ with $\epsilon > 0$ such that $\tau(p) = t$ for all $p \in P$.*

Proof. Suppose, to obtain a contradiction, that the equilibrium involves $\tau(p) = t$ for all $p \in P$. Then, let $p_i \in P$. Firm i's (time t) payoff after the game has reached (without experimentation) time t is

$$v(t \mid p_i) = p_i\left[\left(G(\bar{p}+\epsilon) - G(\bar{p})\right)\Pi_1 + \left(1 - G(\bar{p}+\epsilon) + G(\bar{p})\right)\Pi_2\right],$$

since with probability $G(\bar{p}+\epsilon) - G(\bar{p})$ firm $\sim i$ has $p \in P$ and thus also experiments at time t. In this case, firm i, when successful, is not copied and receives Π_1. With

the complementary probability, it is copied and receives Π_2. Now consider the deviation $\tau(p_i) = t + \delta$ for $\delta > 0$ and arbitrarily small. The payoff to this is

$$v_d\left(t \mid p_i\right) = e^{-r\delta}\left[\left(G\left(\bar{p}+\epsilon\right) - G\left(\bar{p}\right)\right)\left(\zeta\Pi_2 + \left(1-\zeta\right)p_i\Pi_1\right) + \left(1 - G\left(\bar{p}+\epsilon\right) + G\left(p_1\right)\right)p_i\Pi_2\right],$$

where $\zeta \equiv \mathbb{E}\left[p \mid p \in P\right]$ is the expected probability of success of a firm with type in the set P. Since $\Pi_2 > p_i\Pi_1$, we have $\zeta\Pi_2 + (1-\zeta)\,p_i\Pi_1 > p_i\Pi_1$. Moreover, by the assumption that G has strictly positive density, $G\left(\bar{p}+\epsilon\right) - G\left(\bar{p}\right) > 0$. Thus for δ sufficiently small, the deviation is profitable. This contradiction establishes the lemma. ■

Lemma 6. *$\tau(p)$ is continuous in $[a, b]$.*

Proof. Suppose $\tau\left(p\right)$ is discontinuous at $\bar{p}$. Assume without loss of generality that $\tau\left(\bar{p}+\right) \equiv \lim_{\epsilon\downarrow 0}\tau\left(\bar{p}+\epsilon\right) > \tau\left(\bar{p}-\right) \equiv \lim_{\epsilon\uparrow 0}\tau\left(\bar{p}+\epsilon\right)$. Then firms with signal $p = \bar{p}+\delta$ for sufficiently small $\delta > 0$ can experiment at time $\tau\left(\bar{p}-\right) + \epsilon$ for $\epsilon < \tau\left(\bar{p}+\right) - \tau\left(\bar{p}-\right)$ and increase their payoff since $r > 0$. ■

Lemma 7. *$\tau(p)$ is strictly monotone on $[a, b]$.*

Proof. Suppose, to obtain a contradiction, that there exist $q_1 > q_2$ such that $\tau(q_1) = \tau(q_2) = \bar{\tau}$. Suppose that $\sim i$ follows the equilibrium strategy characterized by $\tau(p)$ and consider firm i's expected profit when $p_i = q$ and it chooses to experiment at time t. This can be written as

$$V(q,t) = \int_{p\in P^t_{before}} e^{-r\tau(p)}\left(p\Pi_2 + (1-p)\,q\Pi_1\right)dG(p) + e^{-rt}q\Pi_2\int_{p\in P^t_{after}} dG(p), \qquad \text{(A.2)}$$

where $P^t_{before} = \{p\colon\ \tau(p) \leq t\}$ and $P^t_{after} = \{p\colon\ \tau(p) > t\}$. Notice that $V(q,t)$ is linear in q.

For $\tau(p)$ to characterize a symmetric equilibrium strategy and given our assumption that $\tau(q_1) = \tau(q_2) = \bar{\tau}$, we have

$$V\left(q_1, \bar{\tau}\right) \geq V\left(q_1, t'\right) \text{ and } V\left(q_2, \bar{\tau}\right) \geq V\left(q_2, t'\right) \qquad \text{(A.3)}$$

for all $t' \in \mathbb{R}_+$.

Now take $q = \alpha q_1 + (1-\alpha)q_2$ for some $\alpha \in (0,1)$. By the linearity of $V(q,t)$, this implies that for any $t \neq \bar{\tau}$, we have

$$\begin{aligned} V(\alpha q_1 + (1-\alpha)q_2, t) &= \alpha V(q_1, t) + (1-\alpha) V(q_2, t) \\ &\leq \alpha V(q_1, \tau(q_1)) + (1-\alpha) V(q_2, \tau(q_2)) \\ &= V(\alpha q_1 + (1-\alpha)q_2, \bar{\tau}), \end{aligned}$$

where the middle inequality exploits (A.3). This string of inequalities implies that

$$\tau(\alpha q_1 + (1-\alpha)q_2) = \bar{\tau} \text{ for } \alpha \in [0,1].$$

Therefore, τ must be constant between q_1 and q_2. But this contradicts Lemma 5, establishing the current lemma. ■

The three lemmas together establish that τ is continuous and strictly monotone. This implies that τ is invertible, with inverse $\tau^{-1}(t)$. Moreover, $\tau(b) = 0$, since otherwise a firm with signal $b - \epsilon$ could experiment earlier and increase its payoff. Now consider the maximization problem of firm i with signal q. This can be written as an optimization problem where the firm in question chooses the threshold signal $p = \tau^{-1}(t)$ rather than choosing the time of experimentation t. In particular, this maximization problem can be written as

$$\max_{p \in [a,b]} \int_p^b e^{-r\tau(p_{\sim i})} \left(p_{\sim i}\Pi_2 + (1-p_{\sim i})q\Pi_1\right) dG(p_{\sim i}) + e^{-r\tau(p)} G(p)\, q\Pi_2, \tag{A.4}$$

where the first term is the expected return when the firm $\sim i$ has signal quality $p_{\sim i} \in [p, b]$ and the second term is the expected return when $p_{\sim i} < p$, so that firm $\sim i$ will necessarily copy from i's successful innovation.

Next, suppose that τ is differentiable (we will show below that τ must be differentiable).

Then the objective function (A.4) is also differentiable and the first-order optimality condition can be written (after a slight rearrangement) as

$$r\tau'(p) = \frac{g(p)}{G(p)}\left[1 - \frac{p}{q} - (1-p)\beta^{-1}\right].$$

In a symmetric equilibrium, the function $\tau(p)$ must be a best response to itself, which here corresponds to $p = q$. Therefore, when differentiable, $\tau(p)$ is a solution to

$$r\tau'(p) = -\frac{g(p)}{G(p)}(1-p)\beta^{-1}. \tag{A.5}$$

Integrating this expression, then using integration by parts and the boundary condition $\tau(b) = 0$, we obtain the unique solution (when $\tau(p)$ is differentiable) as

$$\begin{aligned}\tau(p) &= \frac{1}{r\beta}\int_p^b (1-z)\frac{g(z)}{G(z)}dz \\ &= \frac{1}{r\beta}\left[\log G(b)(1-b) - \log G(p)(1-p) + \int_p^b \log G(z)dz\right].\end{aligned}$$

To complete the proof, we need to establish that this is the unique solution. Lemmas 6 and 7 imply that $\tau(p)$ must be continuous and strictly monotone. The result follows if we prove that $\tau(p)$ is also differentiable. Recall that a monotone function is differentiable almost everywhere, i.e., it can can have at most a countable number of points of non-differentiability (see, for example, [28], p. 101, Theorem 3.23). Take $\bar{p}$ to be a point of non-differentiability. Then there exists some sufficiently small $\epsilon > 0$ such that $\tau(p)$ is differentiable on $(\bar{p}-\epsilon, \bar{p})$ and on $(\bar{p}, \bar{p}+\epsilon)$. Then (A.5) holds on both of these intervals. Integrating it over these intervals, we obtain

$$\begin{aligned}\tau(p) &= \tau(\bar{p}-\epsilon) - \frac{1}{r\beta}\int_{\bar{p}-\epsilon}^p (1-z)\frac{g(z)}{G(z)}dz \qquad \text{for } p \in (\bar{p}-\epsilon, \bar{p}), \text{ and} \\ \tau(p) &= \tau(\bar{p}) - \frac{1}{r\beta}\int_{\bar{p}}^p (1-z)\frac{g(z)}{G(z)}dz \qquad \text{for } p \in (\bar{p}, \bar{p}+\epsilon).\end{aligned}$$

Now taking the limit $\epsilon \to 0$ on both intervals, we have either (i) $\tau(\bar{p}+) \neq \tau(\bar{p}-)$; or (ii) $\tau(\bar{p}+) = \tau(\bar{p}-)$. The first of these two possibilities contradicts continuity, so (ii) must apply. But then $\tau(\bar{p})$ is given by (A.1) and is thus differentiable. This argument establishes that $\tau(p)$ is differentiable everywhere and proves the uniqueness of equilibrium. ■

Step 2: To complete the proof of Proposition 12 we need to consider the more general case when the support of G is $[a, b] \subset [0, 1]$. This consists of the showing three additional claims. First, we show that firms with $p \geq \beta$ will always experiment before firms with $p \in [a, \beta)$. The claim follows by a single-crossing argument. First, recall that the value of experimenting at time t for a firm with $p \in [a, \beta)$ is given by (A.2). Defining $P^t_{after\wedge\beta+} = \{p\colon \tau(p) > t \text{ and } p \geq \beta\}$ and $P^t_{after\wedge\beta-} = \{p\colon \tau(p) > t \text{ and } p < \beta\}$, the value of experimenting for a firm with $q \in [a, \beta)$ can be rewritten as

$$V(q,t) = \\ q\Pi_2 \left\{ \int_{p \in P^t_{before}} e^{-r\tau(p)} \left(\frac{p}{q} + \frac{1-p}{\beta} \right) dG(p) + e^{-rt} \left[\frac{1}{\beta} \int_{p \in P^t_{after\wedge\beta+}} dG(p) + \int_{p \in P^t_{after\wedge\beta-}} dG(p) \right] \right\}, \tag{A.6}$$

which exploits the fact that when $p \in P^t_{before}$ or when $p \in P^t_{after\wedge\beta+}$, there will be no copying, and when $p \in P^t_{after\wedge\beta-}$, the innovation (which takes place again with probability q) will be copied, for a payoff of $\Pi_2 = \beta\Pi_1$.

Next, turning to firms with $p = q' \geq \beta$, recall that these firms prefer not to copy prior successful experimentation (from Lemma 3). Therefore, their corresponding value can be written as

$$\tilde{V}(q',t) = q'\Pi_2 \left\{ \frac{1}{\beta} \int_{p \in P^t_{before}} e^{-r\tau(p)} dG(p) + e^{-rt} \left[\frac{1}{\beta} \int_{p \in P^t_{after\wedge\beta+}} dG(p) + \int_{p \in P^t_{after\wedge\beta-}} dG(p) \right] \right\}, \tag{A.7}$$

Note also that when the experimentation time is reduced, say from t to $t' < t$, the first integral gives us the cost of such a change and the second expression (e^{-rt} times the

square bracketed term) gives the gain. Now the comparison of (A.6) to (A.7) establishes the single-crossing property, meaning that at any t a reduction to $t' < t$ is always strictly more valuable for $q' \geq \beta$ than for $q \in [a, \beta)$. First, the gains, given by the expression in (A.6) and (A.7) are identical. Second, the term in parenthesis in the first integral in (A.6) is a convex combination of $1/q > 1/\beta$ and $1/\beta$, and thus is strictly greater than $1/\beta$, so that the cost is always strictly greater for $q \in [a, \beta)$ than for $q' \geq \beta$. From this strict single-crossing argument it follows that there exists some T such that $\tau(p) \leq T$ for all $p \geq \beta$ and $\tau(p) > T$ for all $p \in [a, \beta)$.

The second claim establishes that all firms with $p \geq \beta$ will experiment immediately, that is, $\tau(p) = 0$ for all $p \geq \beta$. To show this, first note that all terms in (A.7) are multiplied by $q' \geq \beta$, so the optimal set of solutions for any firm with $p \geq \beta$ must be identical. Moreover, since $\tau(p) > T$ for all $p \in [a, \beta)$, $P^t_{after \wedge \beta -}$ is identical for all $t \in [0, T]$, and $t > 0$ is costly because $r > 0$. Therefore, the unique optimal strategy for all $p \geq \beta$ is to experiment immediately. Therefore, $\tau(p) = 0$ for all $p \geq \beta$.

Finally, we combine the equilibrium behavior of firms with $p \geq \beta$ with those of $p \in [a, \beta)$. First, suppose that $b \leq \beta$. Then the characterization in Proposition 20 applies exactly. Next suppose that $b > \beta$, so that some firms might have signals $p \geq \beta$. The previous step of the proof has established that these firms will experiment immediately. Subsequently, firms with $p \in [a, \beta)$ will copy a successful innovation at time $t = 0$ or experiment if there is an unsuccessful experimentation at $t = 0$. If there is no experimentation at $t = 0$, then equilibrium behavior (of firms with $p \in [a, \beta)$) is given by Proposition 20 except that the upper support is now β and the relevant distribution is $G(p)$ conditional on $p \in [a, \beta)$, thus all terms are divided by $G(\beta)$. This completes the proof of Proposition 12.

Proof of Proposition 13

The proof mimics that of Proposition 12, with the only difference that the maximization problem of firm i, with signal $p_i = q$, is now modified from (A.4) to

$$\max_{p\in[a,b]} \int_p^b e^{-r\tau(p_{\sim i})} \left(p_{\sim i}\left(\Pi_2 - \eta\right) + \left(1 - p_{\sim i}\right) q\Pi_1\right) dG\left(p_{\sim i}\right) + e^{r\tau(p)} G\left(p\right) q\left(\Pi_2 + \eta\right),$$

which takes into account that copying has cost η and if firm i is the first innovator, then it will be copied and will receive η. Repeating the same argument as in Proposition 12 establishes that the unique equilibrium is given by (4.5).

To prove the second part of the proposition, first suppose that $b < p^{\eta}$ and $b < p^{\eta'}$ so that $\bar{b} = b$ in both cases. Recall also that $\tau^{\eta}(p) = \bar{\tau}^{\eta}(p) > 0$ for $p \in [a, p^{\eta})$. p^{η} is decreasing in η, so that $\tau^{\eta}(p) = 0$ implies that $\tau^{\eta'}(p) = 0$ for any $\eta' > \eta$. We therefore only need to show that $\bar{\tau}^{\eta}(p)$ is strictly decreasing for all $p \in [a, p^{\eta})$. Since $\bar{\tau}^{\eta}(p)$ is differentiable, it is sufficient to show that its derivative with respect to η is negative. This follows by differentiating (4.5) (with $\bar{b} = b$). In particular,

$$\frac{d\tau^{\eta}(p)}{d\eta} = -\frac{1}{(\Pi_2 + \eta)}\left[\bar{\tau}^{\eta}(p) + \frac{2}{rG\left(\bar{b}\right)}\left(\log G\left(\bar{b}\right) - \log G\left(p\right)\right)\right] < 0,$$

since $\log G(\bar{b}) > \log G(p)$ and $\bar{\tau}^{\eta}(p) > 0$.

Next, suppose that $b > p^{\eta'}$. In that case $\bar{b} = p^{\eta'}$ and $d\tau^{\eta}(p)/d\eta$ (in the neighborhood of η') will include additional terms because of the effect of η on $\bar{b}$. In particular:

$$\begin{aligned}\frac{d\tau^{\eta'}(p)}{d\eta} &= -\frac{1}{(\Pi_2 + \eta')}\left[\bar{\tau}^{\eta'}(p) + \frac{2}{rG\left(p^{\eta'}\right)}\left(\log G\left(p^{\eta'}\right) - \log G\left(p\right)\right)\right] \\ &\quad -\frac{g(p^{\eta'})}{G(p^{\eta'})\Pi_1}\frac{\Pi_1 - 2\eta' - p^{\eta'}\Pi_1}{r(\Pi_2 + \eta')G\left(p^{\eta'}\right)} + \frac{g(p^{\eta'})}{G(p^{\eta'})\Pi_1}\bar{\tau}^{\eta'}(p).\end{aligned} \tag{A.8}$$

The first line is again strictly negative and so is the first expression in the second line. The second expression in the second line could be positive, however. For given η', this term

is decreasing in p and tends to 0 as p approaches $p^{\eta'}$ (from (4.5)). Therefore, there exists $p^*(\eta')$ such that for $p \geq p^*(\eta')$, it is no larger than the first term in the second line. This establishes that for $p \in [p^*(\eta'), p^{\eta'})$, $\tau^{\eta'}(p)$ is again decreasing in η, completing the proof.

Appendix B

Equilibrium with Private Signals and Multiple Firms

We now extend the analysis of the private and heterogeneous signals environment from Section 4 to the case with multiple firms. In particular, we show that all equilibria are payoff equivalent for all players and involve firms with strong signals ($p_i \geq \beta$) experiment first. We also discuss briefly the structure of those equilibria and provide an explicit characterization of the unique mixed equilibrium, when all firms with signals $\geq \beta$ use the same experimentation strategy.

Proposition 21. *Suppose that there are $N \geq 3$ firms and the support of G satisfies $[a, b] \not\subset [0, \beta]$. Then:*

1. *There does not exist a symmetric equilibrium in which all firms with $p \geq \beta$ experiment at $t = 0$.*

2. *All symmetric equilibria involve firms with $p \geq \beta$ experimenting in the time interval $[0, T]$ and the rest of the firms experiment after T (if there is no prior experimentation) for some $T > 0$.*

3. *All symmetric equilibria take the following form:*

(a) *Firms with $p \geq \beta$ experiment in time interval $[0, T]$ with flow rate of experimentation $\xi(p, t)$.*

(b) *If there is not any prior experimentation, a firm with signal $p_i < \beta$ experiments at time $\tau(p_i) > T$, where $\tau(.)$ is a strictly decreasing function.*

Moreover, all such equilibria are payoff equivalent for all players.

Note that the characterization of Part 3 allows for pure strategies from firms with strong signals ($p \geq \beta$) - in fact there is such an equilibrium.

Proof. (Part 1) Suppose that $N = 3$, and that $\Pi_2 = \Pi_3$. Suppose, to obtain a contradiction, that there exists a symmetric equilibrium where all firms with $p \geq \beta$ experiment at $t = 0$. Consider firm i with $p_i > \beta$. Let χ_0 be the probability that none of the other two firms have $p \geq \beta$, χ_1 be the probability that one of the other two firms has $p \geq \beta$ and χ_2 be the probability that both firms have $p \geq \beta$. Let us also define $\zeta = \mathbb{E}[p \mid p \geq \beta]$. Since $p_i > \beta$, by hypothesis, firm i experiments at time $t = 0$. Its expected payoff is

$$V(p_i, 0) = \chi_0 p_i \Pi_2 + \chi_1 p_i \left[\zeta \left(\frac{\Pi_1}{2} + \frac{\Pi_2}{2}\right) + (1 - \zeta)\Pi_2\right] + \chi_2 p_i \Pi_1.$$

Intuitively, when none of the other two firms have $p \geq \beta$, when successful, the firm is copied immediately, receiving payoff Π_2. When both of the other two firms have $p \geq \beta$, there is no copying, so when successful, firm i receives Π_1. When one of the other two firms has $p \geq \beta$, then this other firm also experiments at time $t = 0$ and is successful with probability $\zeta = \mathbb{E}[p \mid p \geq \beta]$. In that case, in a symmetric equilibrium the third firm copies each one of the two successful innovations with probability $1/2$. With the complementary probability, $1 - \zeta$, the other firm with $p \geq \beta$ is unsuccessful, and the third firm necessarily copies firm i.

Now consider the deviation to wait a short interval $\epsilon > 0$ before innovation. This will

have payoff

$$\begin{aligned}\lim_{\epsilon\downarrow 0} V\left(p_i,\epsilon\right) &= \chi_0 p_i \Pi_2 + \chi_1 p_i \left(\zeta \Pi_1 + \left(1-\zeta\right)\Pi_2\right) + \chi_2 p_i \Pi_1 \\ &> V\left(p_i,0\right).\end{aligned}$$

The first line of the previous expression follows since, with this deviation, when there is one other firm with $p \geq \beta$, the third firm necessarily will copy the first innovator. The inequality follows since $\Pi_1 > \Pi_2$, establishing that there cannot be an equilibrium in which all firms with $p \geq \beta$ experiment at time $t = 0$. This argument generalizes, with a little modification, to cases in which $N > 3$ and Π_ns differ.

(Part 2) Part 2 follows from a single crossing argument similar to the one used in the proof of Proposition 12.

(Part 3 - Sketch) We give an argument for why all equilibria are payoff equivalent for all players. Consider firms with strong signals, i.e., firms such that $p \geq \beta$. Note that $\xi(p,t) > 0$ for at least some p for all $t \in [0,T]$. It is evident now from (A.7) in the proof of Proposition 20 that all firms with signal $p \geq \beta$ have to be indifferent between experimenting at any time t in $[0,T]$ and, in particular, between experimenting at t and at 0. Combined with Part 1, this implies that the expected payoff of a firm with signal $p \geq \beta$ is equal to

$$p\Pi_1 \mathbb{P}(p_i \geq \beta, \text{ for all } i) + p\Pi_2 \left(1 - \mathbb{P}(p_i \geq \beta, \text{ for all } i)\right),$$

in all equilibria (no matter what the equilibrium strategy profile is). Similarly, we can show that the expected payoff for the rest of the firms (firms with weak signals) is also the same in all equilibria. ■

In Proposition 21 we did not explicitly describe any equilibrium. Proposition 22 provides a characterization of the unique mixed strategy equilibrium, when firms with signals $p \geq \beta$ use the same experimentation strategy, i.e., the rate of experimentation for those

firms depends only on time t (not their signal quality), $\xi(p,t) = \xi(t)$. To simplify the exposition, we focus on an economy in which $N = 3$ and G is uniform on $[0,1]$. The characterization result in this proposition can be (in a relatively straightforward way) extended to $N > 3$. We also conjecture that it can be extended to any distribution G, though this is less trivial.

Proposition 22. *Consider an economy with $N = 3$ firms, $\Pi_2 = \Pi_3$ and G uniform over $[0,1]$. Then, the following characterizes the unique symmetric equilibrium, when firms with signals $p \geq \beta$ use the same experimentation strategy.*

- *Firms with $p \geq \beta$ experiment at the flow rate $\xi(t)$ as long as no other firm has experimented until t. They experiment immediately following another (successful or unsuccessful) experiment. There exists $T < \infty$ such that*

$$e^{-\int_0^T \xi(t)dt} = 0.$$

  That is, all firms with $p \geq \beta$ will have necessarily experimented within the interval $[0,T]$ (or equivalently, $\lim_{t\to T}\xi(t) = +\infty$).

- *Firms with $p < \beta$ immediately copy a successful innovation and experiment at time $\tilde{\tau}_2(p)$ following an unsuccessful experimentation and at time $\tilde{\tau}_3(p)$ if there has been no experimentation until time T.*

Proof. Let us define $\mu(t)$ as the probability that firm $\sim i$ that has not experimented until time t has $p_{\sim i} \geq \beta$. The assumption that G is uniform over $[0,1]$ implies that $\mu(0) = 1 - \beta$.

Now consider the problem of firm i with $p_i \geq \beta$. If there has yet been no experimentation and this firm experiments at time t, its payoff (discounted to time $t = 0$) is

$$V(p_i, t) = p_i e^{-rt}[\Pi_1 \mu(t)^2 + \Pi_2(1 - \mu(t)^2)],$$

since $\mu(t)^2$ is the probability with which both other firms have $p \geq \beta$ and will thus not copy. With the complementary probability, its innovation will be copied. Alternately, if it delays experimentation by some small amount $dt > 0$, then its payoff is:

$$
\begin{aligned}
& V(p_i, t + dt) = \\
& p_i e^{-r(t+dt)} \Big[2\bar{p}\xi(t)\mu(t)dt\Pi_1 \qquad +(1 - 2\xi(t)\mu(t)dt)[\Pi_1\mu(t+dt)^2 + \Pi_2(1 - \mu(t+dt)^2)] \\
& \qquad\qquad\qquad\qquad\qquad\quad +2\xi(t)\mu(t)(1-\bar{p})dt[\Pi_1\mu(t+dt) + \Pi_2(1 - \mu(t+dt))] \Big],
\end{aligned}
$$

where $\bar{p} \equiv \mathbb{E}[p|p > \beta]$ and we use the fact that other firms with $p \geq \beta$ experiment at the rate $\xi(t)$. In a mixed-strategy equilibrium, these two expressions must be equal (as $dt \to 0$). Setting these equal and rearranging, we obtain

$$
\frac{d(\mu^2(t))}{dt} + 2\xi(t)\mu(t)(1-\mu(t))[\bar{p} + \mu(t)] - r\mu^2(t) = \frac{\Pi_2 r}{\Pi_1 - \Pi_2}. \tag{B.1}
$$

In addition, the evolution of beliefs $\mu(t)$ given the uniform distribution and flow rate of experimentation at $\xi(t)$ can be obtained as

$$
\mu(t) = \frac{e^{-\int_0^t \xi(\tau)d\tau}(1-\beta)}{e^{-\int_0^t \xi(\tau)d\tau}(1-\beta) + \beta}. \tag{B.2}
$$

Now let us define

$$
f(t) \equiv e^{-\int_0^t \xi(\tau)d(\tau)}(1-\beta). \tag{B.3}
$$

Using (B.3), (B.2) can be rewritten as

$$
\mu(t) = \frac{f(t)}{f(t) + \beta},
$$

which in turn implies

$$
f(t) = \frac{\beta\mu(t)}{1 - \mu(t)}.
$$

Moreover (B.3) also implies that

$$\xi(t) = -\frac{f'(t)}{f(t)} = -\frac{\mu'(t)}{(1-\mu(t))\mu(t)}. \tag{B.4}$$

Substituting these into (B.1), we obtain the following differential equation for the evolution of $\mu(t)$:

$$2\mu(t)\mu'(t) + 2\xi(t)\mu(t)(1-\mu(t))[\bar{p}+\mu(t)] - r\mu^2(t) = \frac{r}{\beta^{-1}-1}.$$

Further substituting $\xi(t)$ from (B.4), we obtain

$$\mu'(t) = -\frac{r}{2\bar{p}(\beta^{-1}-1)}(1+\left(\beta^{-1}-1\right)\mu^2(t)).$$

This differential equation satisfies the Lipschitz condition and therefore it has a unique solution, which takes the form

$$\mu(t) = \frac{1}{\sqrt{\beta^{-1}-1}}\tan\left(\sqrt{\beta^{-1}-1}\left[-\frac{r}{2\bar{p}\left(\beta^{-1}-1\right)}t + \frac{\arctan\left(\mu(0)\sqrt{\beta^{-1}-1}\right)}{\sqrt{\beta^{-1}-1}}\right]\right),$$

with boundary condition $\mu(0) = 1-\beta$. Given this solution, the flow rate of experimentation for firms with $p \geq \beta$, $\xi\left(t\right)$, is obtained from (B.4) as

$$\xi(t) = c_1c_2(1+\tan(c_1(-c_2t+c_3))^2)\left[-\frac{1}{-c_1+\tan(c_1(-c_2t+c_3))} + \frac{1}{\tan(c_1(-c_2t+c_3))}\right],$$

where

$$c_1 \equiv \sqrt{\beta^{-1}-1},\ c_2 \equiv -\frac{r}{2\bar{p}\left(\beta^{-1}-1\right)},\ \text{and } c_3 \equiv \frac{\arctan\left(\mu(0)\sqrt{\beta^{-1}-1}\right)}{\sqrt{\beta^{-1}-1}}.$$

It can then be verified that

$$\lim_{t\to T}\xi\left(t\right) = \infty,$$

where $T = c_3/c_2$. It can also be verified that for all $t \in [0, T]$, where firms with $p \geq \beta$ are experimenting at positive flow rates, firms with $p < \beta$ strictly prefer to wait. The equilibrium behavior of these firms after an unsuccessful experimentation or after time T is reached is given by an analysis analogous to Proposition 12. Combining these observations gives the form of the equilibrium described in the proposition. ■

Appendix C

Experimentation in Discrete Time

We discuss a discrete time version of the model described in the main text. We formally show that the continuous time model, which is our main focus, provides the same economic and mathematical answers as the limit of the discrete-time model, when the length of the time interval $\Delta \to 0$. We limit the discussion to the case of two symmetric firms.

Environment

Let us denote the time interval between two consecutive periods by $\Delta > 0$. In what follows we will take Δ to be small. During an interval of length Δ, the payoff to a firm that is the only one implementing a successful project is $\pi_1 \Delta > 0$. In contrast, if a successful project is implemented by both firms, each receives $\pi_2 \Delta > 0$. The payoff to an unsuccessful project is normalized to zero. Both firms discount the future at the common rate $r > 0$ (so that the discount factor per period is $e^{-r\Delta}$).

Strategies are defined in this game as follows. Let a history up to time t (where $t = k\Delta$ for some integer k) be denoted by h^t. The set of histories is denoted by $\mathcal{H}^t$. A strategy for a firm is a mapping from the history of the game up to time t, h^t, to the probability of experimentation at a given time interval and the set of projects. Thus the time t strategy can be written as

$$\sigma^t : \mathcal{H}^t \to [0,1] \times \{1,2\},$$

where $[0,1]$ denotes the probability of implementing a project (either experimenting or copying) at time t.

Equilibria

We start with asymmetric equilibria. In an asymmetric equilibrium, one of the firms, say 1, immediately experiments with its project. Firm 2 copies firm 1 in the next time period if the latter is successful and tries its own project otherwise. In terms of the notation above, this asymmetric equilibrium would involve

$$\hat{\sigma}_1^t(\cdot) = (1,1),$$

for $t = 0, \Delta, 2\Delta, \ldots$. In words, this means that firm 1 chooses to experiment immediately (if it has not experimented yet until t) and experiments with its project. Firm 2, on the other hand, uses the strategy

$$\hat{\sigma}_2^t\left(a^t, s^t\right) = \begin{cases} (1,1) & \text{if } a^t = 1 \text{ and } s^t = 1, \\ (1,2) & \text{if } a^t = 1 \text{ and } s^t = 0, \\ (0,\cdot) & \text{if } a^t = 0, \end{cases}$$

for $t = 0, \Delta, 2\Delta, \ldots$. The crucial feature highlighted by these strategies is that firm 2 never experiments until firm 1 does.

Using the same analysis as in the proof of Proposition 2 below, it is straightforward to verify that $\hat{\sigma}_2$ is a best response to $\hat{\sigma}_1$ provided that $\Delta < \Delta^* \equiv r^{-1}\log(\beta + 1 - p)$. What about $\hat{\sigma}_1$? Given $\hat{\sigma}_2$, suppose that the game has reached time t (where $t = k\Delta$ for $k \in \mathbb{N}$). If firm 1 now follows $\hat{\sigma}_1$, it will receive expected payoff

$$V[t] = p\left(e^{-rt}\pi_1\Delta + e^{-r(t+\Delta)}\Pi_2\right),$$

at time t, since its experimentation will be successful with probability p, yielding a profit of

$\pi_1\Delta$ during the first period following the success (equivalent to $e^{-rt}\pi_1\Delta$ when discounted to time $t = 0$). Then according to $\hat{\sigma}_2$, firm 2 will copy the successful project and firm 1 will receive the present discounted value $e^{-r(t+\Delta)}\Pi_2$ from then on. If, at this point, firm 1 chooses not to experiment, then the game proceeds to time $t + \Delta$, and according to the strategy profile $\left(\hat{\sigma}_1^{t+\Delta}, \hat{\sigma}_2^{t+\Delta}\right)$, it will receive payoff equal to

$$V[t+\Delta] = e^{-r\Delta}V[t] < V[t].$$

Therefore this deviation is not profitable. This discussion establishes the following proposition (proof in the text).

Proposition 23. *Suppose that Assumption 1 holds and that $\Delta < \Delta^* \equiv r^{-1}\log(\beta + 1 - p)$. Then there exist two asymmetric equilibria. In each, one firm, $i = 1, 2$, tries its project with probability 1 immediately and the other, firm $\sim i$, never tries its project unless it observes the outcomes of the experimentation of firm i. Following experimentation by i, firm $\sim i$ copies it if successful and experiments with its own project otherwise.*

More formally, the two equilibria involve strategies of the form:

$$\sigma_i^t(\cdot) = (1, i),$$

$$\sigma_{\sim i}^t\left(a^t, s^t\right) = \begin{cases} (1, i) & \text{if } a^t = 1 \text{ and } s^t = 1, \\ (1, \sim i) & \text{if } a^t = 1 \text{ and } s^t = 0, \\ (0, \cdot) & \text{if } a^t = 0, \end{cases}$$

for $i = 1, 2$.

Next we study symmetric equilibria. We focus on the case where the time interval Δ is strictly positive but small.

As defined above a firm's strategy is a mapping from its information set to the probability of implementing a project. We refer to a strategy as pure if the experimentation

probability at a given time t is either 0 or 1. That is, a pure strategy takes the form

$$\sigma^t : \mathcal{H}^t \to \{0, 1\} \times \{1, 2\}.$$

Our first result shows that for small Δ, there are no pure-strategy symmetric equilibria.

Proposition 24. *Suppose that Assumption 1 holds and that $\Delta < \Delta^* \equiv r^{-1} \log(\beta + 1 - p)$ (where recall that $\beta \equiv \Pi_2/\Pi_1$). Then there exist no symmetric pure-strategy equilibria.*

Proof. Suppose, to obtain a contradiction, that such an equilibrium σ^* exists. The first case we need to consider, is when σ^* involves no experimentation, i.e., both firms wait with probability 1 for every time t. Then, it is straightforward to see, that firm 1 experimenting with probability 1 at time $t = 0$, is a profitable deviation. This implies that for such an equilibrium σ^* to exist, there should exist some time t_0 and history h^{t_0} (with $t_0 = k\Delta$ for $k \in \mathbb{N}$) such that $\sigma^{t_0}(\varphi = j, h^t) = (1, j)$. Then following this history the payoff to both firms is

$$V[t_0 \mid \sigma^*, \sigma^*] = e^{-rt_0} p\Pi_1.$$

Now consider a deviation by firm 1 to σ', which involves, after history h^{t_0}, waiting until date $t_0 + \Delta$, copying firm 2's project if successful, and experimenting with its own project otherwise. The payoff to this strategy is

$$V[t_0 \mid \sigma', \sigma^*] = e^{-r(t_0+\Delta)} (p\Pi_2 + (1-p)p\Pi_1),$$

since firm 2 experiments with probability 1 at time t_0 and is successful with probability p. Clearly, $V[t_0 \mid \sigma', \sigma^*] > V[t_0 \mid \sigma^*, \sigma^*]$ and there is a profitable deviation if

$$e^{-r(t_0+\Delta)} (p\Pi_2 + (1-p)p\Pi_1) > e^{-rt_0} p\Pi_1,$$

or if

$$\Delta < \Delta^* \equiv r^{-1} \log (\beta + 1 - p) .$$

Here $\Delta^* > 0$ since, from Assumption 1, $\beta \equiv \Pi_2/\Pi_1 > p$. This establishes the existence of a profitable deviation and proves the proposition. ■

Proposition 24 also implies that all symmetric equilibria must involve mixed strategies. Moreover, any candidate equilibrium strategy must involve copying of a successful project in view of Assumption 1 and immediate experimentation when the other firm has experimented. Therefore, we can restrict attention to time t strategies of the form

$$\hat{\sigma}_i^t \left(a^t, s^t\right) = \begin{cases} (1, i) & \text{if } a^t = 1 \text{ and } s^t = 1, \\ (1, \sim i) & \text{if } a^t = 1 \text{ and } s^t = 0, \\ (q(t)\Delta, i) & \text{if } a^t = 0, \end{cases} \tag{C.1}$$

for $t = 0, \Delta, 2\Delta, ...$, where $q(t)\Delta$ is the probability of experimenting at time t conditional on no experimentation by either firm up to time t (all such histories are identical, hence we write $q(t)$ instead of $q(h^t)$). Clearly, feasibility requires that $q(t)\Delta \leq 1$.

Next we derive an explicit characterization of the unique symmetric equilibrium as $\Delta \to 0$. In the text, we assume that firms use a constant probability of experimentation over time, i.e., $q(t) = q$ for all t (Proposition 25 relaxes this assumption and establishes uniqueness more generally). We consider a symmetric mixed-strategy equilibrium σ^* and suppose that the game has reached time t without experimentation. Let $v_w [t \mid \Delta]$ and $v_e [t \mid \Delta]$ denote the time t continuation payoffs to firm i when firm $\sim i$ plays σ^* and firm i chooses to wait or to experiment (and period length is Δ).[1] For a mixed-strategy equilibrium to exist, we need

$$v_w [t \mid \Delta] = v_e [t \mid \Delta] . \tag{C.2}$$

The proof of Proposition 25 below shows that all symmetric equilibria involve mixing after

[1]Here we use v, since V denotes the value discounted back to $t = 0$.

any history h^t (with no experimentation up to t), i.e., equation (C.2) holds for all such h^t. Therefore, it suffices to characterize σ^* such that (C.2) holds. First, consider firm i's payoffs from experimenting:

$$v_e[t \mid \Delta] = q\Delta p\Pi_1 + (1 - q\Delta)(p\pi_1\Delta + e^{-r\Delta}p\Pi_2), \tag{C.3}$$

since in this case firm i is successful with probability p and receives continuation value Π_1 if firm $\sim i$ has also experimented during the same time interval (probability $q\Delta$), and it receives $\pi_1\Delta + e^{-r\Delta}\Pi_2$ otherwise (payoff for current time interval plus continuation value). The latter event occurs with probability $1 - q\Delta$.

Similarly, its payoff from waiting is

$$v_w[t \mid \Delta] = e^{-r\Delta}\Big(q\Delta\,(p\Pi_2 + (1-p)p\Pi_1) + (1 - q\Delta)v_w[t + \Delta \mid \Delta]\Big), \tag{C.4}$$

where firm i receives no payoff today and with probability $q\Delta$, firm $\sim i$ experiments, in which case firm i copies if the experimentation is successful and experiments with its own project otherwise, with expected continuation return $p\Pi_2 + (1-p)p\Pi_1$. With probability $1 - q\Delta$, firm $\sim i$ does not experiment, and firm i then receives $v_w[t + \Delta \mid \Delta]$. Adding and subtracting $v_w[t + \Delta \mid \Delta]$ from the left-hand side of (C.4) and rearranging, we obtain

$$v_w[t + \Delta \mid \Delta]\left(1 - (1 - q\Delta)\,e^{-r\Delta}\right) - (v_w[t + \Delta \mid \Delta] - v_w[t \mid \Delta]) = e^{-r\Delta}q\Delta\,(p\Pi_2 + (1-p)\,p\Pi_1)\,.$$

Dividing both sides by Δ and taking the limit as $\Delta \to 0$ yields

$$\lim_{\Delta \to 0} v_w[t + \Delta \mid \Delta] - \frac{1}{r+q}\lim_{\Delta \to 0}\left(\frac{v_w[t + \Delta \mid \Delta] - v_w[t \mid \Delta]}{\Delta}\right) = \frac{q}{r+q}\left[p\Pi_2 + (1-p)\,p\Pi_1\right]. \tag{C.5}$$

From equation (C.3), we see that $v_e[t \mid \Delta]$ does not depend on t. Since equation (C.2) holds for all h^t (thus for all t), we have $v_w[t \mid \Delta] = v_e[t \mid \Delta]$ and $v_w[t + \Delta \mid \Delta] =$

$v_e\left[t+\Delta \mid \Delta\right]$, implying that $v_w\left[t+\Delta \mid \Delta\right] = v_w\left[t \mid \Delta\right]$. Therefore, the second term on the left-hand side of (C.5) must be equal to zero. Moreover, taking the limit as $\Delta \to 0$ in (C.3), we obtain

$$\lim_{\Delta\to 0} v_w\left[t+\Delta \mid \Delta\right] = \lim_{\Delta\to 0} v_e\left[t \mid \Delta\right] = p\Pi_2.$$

Combined with (C.5), this yields

$$q(t) = q^* \equiv \frac{r\beta}{1-p} \text{ for all } t, \tag{C.6}$$

where recall, from (2.1), that $\beta \equiv \Pi_2/\Pi_1$.

The next proposition relaxes the assumption that $q(t)$ is constant for all t and shows that this is indeed the unique symmetric equilibrium.

Proposition 25. *Suppose that Assumption 1 holds and $\Delta \to 0$. Then there exists a unique symmetric equilibrium. In this equilibrium, both firms use the mixed strategy $\hat{\sigma}$ as given in (2.2) with $q(t) = q^*$ as in (C.6).*

Proof. We first show that any symmetric equilibrium must involve mixing after any history $h^t \in \mathcal{H}^t$ along which there has been no experimentation. The argument in the proof of Proposition 24 establishes that after any such history h^t, there cannot be experimentation with probability 1. We next show that there is positive probability of experimentation at time $t = 0$. First note that the equilibrium-path value to a firm, V^*, (discounted back to time $t = 0$), satisfies $V^* \geq p\Pi_2$, since each firm can guarantee this by experimenting at time $t = 0$. This implies that in any equilibrium there must exist some time T such that after time T there is positive probability of experimentation and innovation. Now to obtain a contradiction, suppose that $T > 0$. By the argument preceding the proposition, $\lim_{\Delta\to 0} V_e\left[T \mid \Delta\right] = e^{-rT}p\Pi_2$, and therefore, in any mixed-strategy equilibrium, $V^*\left[T \mid \Delta\right] \to e^{-rT}p\Pi_2$. However, for $T > 0$ this is strictly less than $V^* \geq p\Pi_2$, yielding a contradiction and establishing the desired result. The same argument also establishes

that there cannot exist any time interval (T, T'), with $T' > T$, along which there is no mixing.

Hence, along any history h^t where there has not been an experimentation, both firms must be indifferent between waiting and experimenting. This implies that (C.2) must hold for all t. Let $q(t)\Delta$ denote the probability of experimentation at time t. Firm i's payoff for experimenting at time t is given by an expression similar to equation (C.3),

$$v_e [t \mid \Delta] = q(t)\Delta p\Pi_1 + (1 - q(t)\Delta)p\Pi_2. \tag{C.7}$$

We next show that the probability of experimentation $q(t)$ in a symmetric equilibrium is a continuous function of t. Suppose that $q(t)$ is not continuous at some $\bar{t} \geq 0$. If $q(\bar{t}) < q(\bar{t}+)$ (where $q(\bar{t}+) \equiv \lim_{t \downarrow \bar{t}} q(t)$), it then follows from (C.7) and Assumption 1 that $v_e [\bar{t} \mid \Delta] < v_e [\bar{t}+ \mid \Delta]$. This implies that firm i has an incentive to delay experimentation at time $\bar{t}$. But this contradicts the fact that the symmetric equilibrium must involve mixing at all such t. Similarly, if $q(\bar{t}) > q(\bar{t}+)$, we have $v_e [\bar{t} \mid \Delta] > v_e [\bar{t}+ \mid \Delta]$, implying that firm i will experiment with probability 1 at time $\bar{t}$, again yielding a contradiction. This establishes that $q(t)$ is a continuous function of t.

A derivation similar to that preceding the proposition then shows that equation (C.5) holds when q is replaced by $q(t)$. In particular,

$$\lim_{\Delta \to 0} v_w [t + \Delta \mid \Delta] - \frac{1}{r + q(t)} \lim_{\Delta \to 0} \left(\frac{v_w [t + \Delta \mid \Delta] - v_w [t \mid \Delta]}{\Delta} \right) = \frac{q(t)}{r + q(t)} [p\Pi_2 + (1 - p) p\Pi_1]. \tag{C.8}$$

Since $v_w [t \mid \Delta] = v_e [t \mid \Delta]$ and $v_w [t + \Delta \mid \Delta] = v_e [t + \Delta \mid \Delta]$, we can write the second

term on the left-hand side of equation (C.8) as

$$\begin{aligned}
\lim_{\Delta \to 0} \left(\frac{v_w\left[t+\Delta \mid \Delta\right] - v_w\left[t \mid \Delta\right]}{\Delta} \right) &= \lim_{\Delta \to 0} \left(\frac{v_e\left[t+\Delta \mid \Delta\right] - v_e\left[t \mid \Delta\right]}{\Delta} \right) \\
&= \lim_{\Delta \to 0} \Big(q(t+\Delta) - q(t) \Big) \Delta p (\Pi_1 - \Pi_2) \\
&= 0,
\end{aligned}$$

where the second equality follows from equation (C.7) and the third equality holds by the continuity of the experimentation probability $q(t)$. Substituting for

$$\lim_{\Delta \to 0} v_w\left[t+\Delta \mid \Delta\right] = \lim_{\Delta \to 0} v_e\left[t \mid \Delta\right] = p\Pi_2,$$

in equation (C.8) and solving for $q(t)$ yields $q(t) = q^*$ as in (C.6), completing the proof. ∎

Note that the limit when $\Delta \to 0$, which Proposition 25 shows is well behaved, also corresponds to the (symmetric) equilibrium of the same model set up directly in continuous time, thus showing formally the equivalence of the two.

Appendix D

Omitted Proofs from Chapter 6

Algorithm for observation radius

The following algorithm computes the perfect observation radius for agent i and communication network topology G^n. Specifically, the algorithm computes the sequence of optimal decisions of agent i for any realization of the private signals under the assumption that all agents except i never exit and keep transmitting new information. Assume for simplicity that G^n is connected (otherwise apply the algorithm to the connected component in which agent i resides). Let $t_{end,i}$ denote the maximum distance between i and an agent in G^n. Note that this implies $B^n_{t_{end,i}} = N^n$. The state at communication step t, $0 \leq t \leq t_{end,i}$, is simply the number of private signals the agent has observed thus far and the empirical mean for θ, k_t and θ_t respectively. The algorithm computes the optimal decision for agent i starting from the last possible communication step $t_{end,i}$ and working its way through the beginning ($t = 0$). In particular (we drop subscript i),

Algorithm 1.

(1)

$$OPTdec(t_{end}, k_{t_{end}}, \theta_{t_{end}}) = \theta_{t_{end}}.$$

and $Pay(t_{end}, k_{t_{end}}, \theta_{t_{end}}) = \pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k_{t_{end}}}$

(2) For $t = t_{end} - 1$ *to 0 do:*

(i) Preliminaries:

- *Number of new observations at step t:* $\ell = |B^n_{i,t+1}| - |B^n_{i,t}|$.
- $Pay^{exit}(t, k_t, \theta_t) = \pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k_t}$.
- $Pay^{wait}(t, k_t, \theta_t) = Pay(t+1, k_t + \ell, \theta'_t)$.

(ii) If $Pay^{wait} \leq (1 - r/\lambda) Pay^{exit}$ *then*

$$OPTdec(t, k_t, \theta_t) = \theta_t \text{ and } Pay(t, k_t, \theta_t) = \pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 k^n_t}$$

$$\text{Else } OPTdec(t, k_t, \theta_t) = \text{"wait" and } Pay(t, k_t, \theta_t) = Pay^{wait}(t, k_t, \theta_t).$$

Correctness: Note that if agent i has not exited till the last time period, then, since there is no more observations to be made in the future, she will exit and choose the action that maximizes her expected payoff (see Lemma 4). For any other time period t agent i's expected payoff of exiting taking an action at the current time period is given by Pay^{exit}. Finally, the agent decides to "wait" only if this action leads to higher expected payoff. It is straightforward to see that the computational complexity of the algorithm is $O(n^2)$. To relate the algorithm with computing the perfect observation radius, note that the perfect observation radius coincides with the exit time computed by the algorithm.

Finally, Lemma 8 states that the probability of choosing an action, that is more than ϵ away from the optimal, for agent i, such that $i \in V_k^{n,\sigma}$, i.e., $\mathbb{P}_\sigma\left(M_i^{n,\epsilon} = 0\right)$ is uniformly bounded away from 0.

Lemma 8. *Let $k > 0$ be a constant, such that the k-radius set $V_k^{n,\sigma}$ is non-empty. Then,*

$$\mathbb{P}_\sigma(M_i^{n,\epsilon} = 0) \geq erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k}}\right) \text{ for all } i \in V^n_{k,\sigma},$$

where $erf(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$ *is the error function.*

Proof. The lemma follows from the observation that because of our normality assumption the empirical mean $\hat{\theta}$ after observing ℓ private signals is normally distributed around θ with variance $\rho_{\hat{\theta}}^2 = \bar{\rho}^2/\ell$. Then, the probability that $M_i^{n,\epsilon} = 0$ is simply equal to the probability that a random variable $X \sim N(0, \bar{\rho}^2/\ell)$ does not belong to interval $[-\epsilon, \epsilon]$, i.e.,

$$\mathbb{P}_\sigma(M_i^{n,\epsilon} = 0) = erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/\ell}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/\ell}}\right).$$

The lemma follows since agent $i \in V_k^{n,\sigma}$, thus she takes an irreversible action after observing at most k private signals. ■

Proof of Proposition 15. First, we show that learning fails if condition (6.3) holds, i.e., there exists an $\eta > 0$, such that

$$\limsup_{n\to\infty} \frac{1}{n} \cdot \left|V_k^{n,\sigma}\right| > \eta. \tag{D.1}$$

and $k < \bar{k}$ given by Condition (6.3). From condition (D.1) we obtain that there exists an infinite index set J such that

$$\left|V_k^{n_j}\right| \geq \eta \cdot n_j \text{ for } j \in J.$$

Now restrict attention to index set J, i.e., consider $n = n_j$ for some $j \in J$. Then,

$$\begin{aligned} \mathbb{P}_\sigma\left(\tfrac{1}{n}\textstyle\sum_{i=1}^n M_i^{n,\epsilon} > 1-\epsilon\right) &= \mathbb{P}_\sigma\left(\tfrac{1}{n}\left[\textstyle\sum_{i\in V_k^{n,\sigma}} M_i^{n,\epsilon} + \sum_{i\notin V_k^{n,\sigma}} M_i^{n,\epsilon}\right] > 1-\epsilon\right) \\ &\leq \mathbb{P}_\sigma\left(\tfrac{1}{n}\left[\textstyle\sum_{i\in V_k^{n,\sigma}} M_i^{n,\epsilon} + n - \left|V_k^{n,\sigma}\right|\right] > 1-\epsilon\right) \\ &= \mathbb{P}_\sigma\left(\tfrac{1}{n}\textstyle\sum_{i\in V_k^{n,\sigma}} M_i^{n,\epsilon} > \frac{\left|V_k^{n,\sigma}\right|}{n} - \epsilon\right) \end{aligned},$$

where the inequality follows since we let $M_i^{n,\epsilon} = 1$ for all $i \notin V_k^{n,\sigma}$. Next we use Markov's

inequality to obtain

$$\mathbb{P}_\sigma\left(\frac{1}{n}\sum_{i\in V_k^{n,\sigma}} M_i^{n,\epsilon} > \frac{\left|V_k^{n,\sigma}\right|}{n} - \epsilon\right) \leq \frac{\mathbb{E}_\sigma\left[\sum_{i\in V_k^{n,\sigma}} M_i^{n,\epsilon}\right]}{n\cdot\left(\left|V_k^{n,\sigma}\right|/n-\epsilon\right)}.$$

We can view each summand above as an independent Bernoulli variable with success probability bounded above by $erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k}}\right)$ from Lemma 8. Thus,

$$\begin{aligned}\frac{\mathbb{E}_\sigma\left[\sum_{i\in V_k^{n,\sigma}} M_i^{n,\epsilon}\right]}{n\cdot\left(\left|V_k^{n,\sigma}\right|/n-\epsilon\right)} &\leq \frac{\left|V_k^{n,\sigma}\right|\left(erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k}}\right)-erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k}}\right)\right)}{n\cdot\left(\left|V_k^{n,\sigma}\right|/n-\epsilon\right)} \\ &\leq \frac{\eta}{\eta-\epsilon}\left(erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k}}\right)\right) < 1-\delta,\end{aligned}$$

where the second inequality follows from the fact that n was chosen such that $\left|V_k^{n,\sigma}\right| > \eta\cdot n$. Finally, the last expression follows from $k < \bar{k}$ (cf. Condition (6.3)). We obtain that for all $j \in J$ it holds that

$$\mathbb{P}_\sigma\left(\left[\frac{1}{n_j}\sum_{i=1}^{n_j}(1-M_i^{n_j,\epsilon})\right] > \epsilon\right) \geq \delta.$$

Since J is an infinite index set we conclude that

$$\limsup_{n\to\infty}\mathbb{P}_\sigma\left(\left[\frac{1}{n}\sum_{i=1}^{n}(1-M_i^{n,\epsilon})\right] > \epsilon\right) \geq \delta,$$

thus ϵ,δ-asymptotic learning is incomplete when (6.3) holds.

Next, we prove that Condition (6.2) is sufficient for ϵ,δ-asymptotic learning. As mentioned above, if agent i takes an irreversible action after observing ℓ signals, then the probability that $M_i^{n,\epsilon} = 1$ is equal to the probability that a random variable $X \sim N(0,\bar{\rho}^2/\ell)$ takes a value in interval $[-\epsilon,\epsilon]$. Thus,

$$\mathbb{P}_\sigma(M_i^{n,\epsilon}=1) = erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/\ell}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/\ell}}\right). \tag{D.2}$$

Similarly with above, we have

$$\begin{aligned}\mathbb{P}_\sigma\left(\left[\frac{1}{n}\sum_{i=1}^{n}(1-M_i^{n,\epsilon})\right]>\epsilon\right) &\le \mathbb{P}_\sigma\left(\left[\frac{1}{n}\sum_{i\notin V}(1-M_i^{n,\epsilon})\right]>\epsilon-\frac{|V|}{n}\right)\\ &\le \frac{\mathbb{E}_\sigma\left[\sum_{i\notin V}(1-M_i^{n,\epsilon})\right]}{n\left(\epsilon-|V|/n\right)},\end{aligned}\tag{D.3}$$

where $V=\left\{i \mid \quad |B^n_{i,\tau_i^{n,\sigma}}|\le\hat{k}\right\}$ and the second inequality follows from Markov's inequality. By combining Eqs. (D.2) and (D.3) and letting $k_i^{n,\sigma}$ denote the number of private signals that agent i observed before taking an action,

$$\frac{\mathbb{E}_\sigma\left[\sum_{i\notin V}(1-\tilde{M}_i^n)\right]}{n\left(\epsilon-|V|/n\right)}\le\frac{\sum_{i\notin V}1-\left(erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k_i^{n,\sigma}}}\right)-erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k_i^{n,\sigma}}}\right)\right)}{n\left(\epsilon-|V|/n\right)}.\tag{D.4}$$

We have

$$erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k_i^{n,\sigma}}}\right)-erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2/k_i^{n,\sigma}}}\right)>1-\frac{\delta(\epsilon-\zeta)}{1-\zeta},\tag{D.5}$$

for all $i\notin V$ from the definition of $\hat{k}$ (cf. Condition (6.2)). Thus, combining Eqs. (D.3),(D.4) and (D.5), we obtain

$$\mathbb{P}_\sigma\left(\left[\frac{1}{n}\sum_{i=1}^{n}(1-M_i^{n,\epsilon})\right]>\epsilon\right)<\delta \text{ for all } n>N,$$

where N is a sufficiently large constant, which implies that condition (6.2) is sufficient for asymptotic learning. ■

Proof of Proposition 16. We show that although an agent has potentially access to less information under Assumption 6, perfect asymptotic learning occurs whenever perfect asymptotic learning occurs under Assumption 5.

Proposition 26. *If perfect asymptotic learning occurs in society $\{G^n\}_{n=1}^{\infty}$ under Assumption 5, then perfect asymptotic learning occurs under Assumption 6 along any equilibrium*

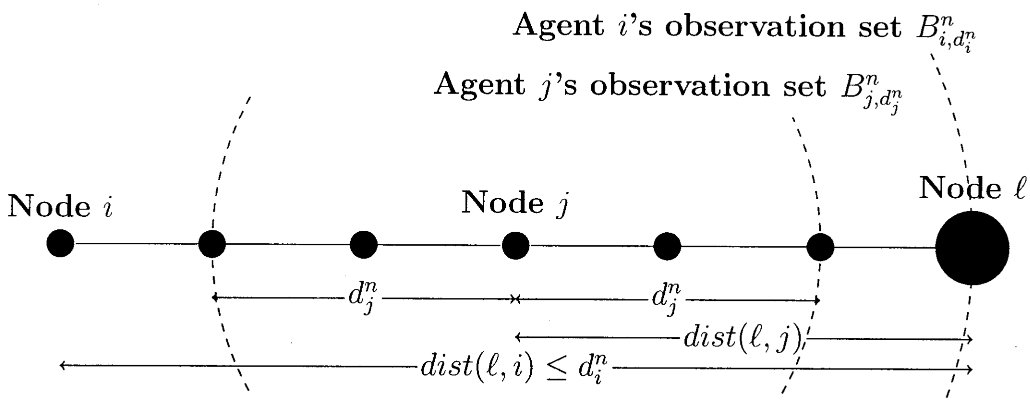

Figure D-1: Proof of Proposition 26.

σ.

Proof. Consider set U_k^n, where recall that U_k^n is the set of agents that are at a short distance (at most equal to their observation radius) from an agent with in-degree at least k. Define similarly set $Z_{k,\sigma}^n$ as the set of agents, that at some equilibrium σ, communicate with an agent with in-degree at least k. Note that under Assumption 5 the sets are equal, i.e., $U_k^n = Z_{k,\sigma}^n$. To complete the proof, we show that for k large enough (and consequently n large enough), $U_k^n = Z_{k,\sigma}^n$, even under Assumption 6.

Consider $i \in U_k^n$ and let $\mathcal{P}^n = \{\ell, i_1, \cdots, i_K, i\}$ denote the shortest path in communication network G^n between i and any agent ℓ, with $deg_\ell^n \geq k$. First we show the following (refer to Figure D-1)

$$i \in U_k^n \Rightarrow j \in U_k^n \text{ for all } j \in \mathcal{P}^n. \tag{D.6}$$

Assume for the sake of contradiction that condition (D.6) does not hold and consider the simplified environment under Assumption 5 (we are only looking at the set U_k^n, so we can restrict attention to that environment). Then, let

$$j = \arg\min_{j'}\{dist^n(\ell, j') \big| j' \in \mathcal{P}^n \text{ and } dist^n(\ell, j') > \tau_{j'}^n\}.$$

For agents i, j we have $\tau_i^n > \tau_j^n$ and $dist(j,i) + d_j^n < dist(\ell, i) \leq \tau_i^n$, since otherwise

$j \in U_k^n$. This implies that $B^n_{j,\tau_j^n} \subset B^n_{i,\tau_i^n}$. Furthermore,

$$-\frac{1}{1/\rho^2 + 1/\bar{\rho}^2(1+|B^n_{j,\tau_j^n}|)} > \mathbb{E}_\sigma(-e^{rt}\frac{1}{1/\rho^2 + 1/\bar{\rho}^2(1+|B^n_{j,\tau_j^n}|+k)}), \tag{D.7}$$

where t is the time for $dist(\ell, j) - \tau_j^n$ extra communication steps to take place. In particular, the left hand side is equal to the expected payoff of agent j if she takes an irreversible action at time τ_j^n after receiving $|B^n_{j,\tau_j^n}|$ observations, whereas the right hand side is a lower bound on the expected payoff if agent j delays taking an action until after she communicates with agent ℓ. The inequality follows, from the definition of the observation radius for agent j. On the other hand,

$$-\frac{1}{1/\rho^2 + 1/\bar{\rho}^2(1+|B^n_{j,\tau_j^n}|)} < \mathbb{E}_\sigma\left(-e^{rt'}\frac{1}{1/\rho^2 + 1/\bar{\rho}^2(1+|B^n_{j,\tau_j^n}|+k)}\right), \text{ for some } \epsilon > 0, \tag{D.8}$$

For k large enough we conclude that $dist(\ell, j) < dist(\ell, i) - dist(j, i)$, which is obviously a contradiction. This implies that (D.6) holds.

Next we show, by induction on the distance from agent ℓ with degree $\geq k$, that $U_k^n = Z^n_{k,\sigma}$ for any equilibrium σ. The claim is obviously true for all agents with distance equal to 0 (agent ℓ) and 1 (her neighbors). Assume that the claim holds for all agents with distance at most t from agent ℓ, i.e., if $i \in U_k^n$ and $dist(\ell, i) \leq t$ then $i \in Z^n_{k,\sigma}$. Finally, we show the claim for an agent i such that $i \in U_k^n$ and $dist(\ell, i) = t + 1$. Consider a shortest path $\mathcal{P}^n$ from i to ℓ. Condition (D.6) implies that all agents j in the shortest path are such that $j \in U_k^n$, thus from the induction hypothesis we obtain $j \in Z^n_{k,\sigma}$. Thus, for k sufficiently large we obtain that $i \in Z^n_{k,\sigma}$, for any equilibrium σ.

Finally, from Corollary 2 we conclude that asymptotic learning under Assumption 5 implies asymptotic learning under Assumption 6. ■

The first part of Proposition 16 follows directly from Proposition 15 and Proposition 26. To conclude the proof of Proposition 16 we need to show that if asymptotic learning occurs

when condition (6.2) does not hold along some equilibrium σ, then the society contains a set of leaders. In particular, consider a society $\{G^n\}_{n=1}^{\infty}$ in which condition (6.2) does not hold, i.e., assume as in the proof of Proposition 15 that there exists a $k > 0$, $\epsilon > 0$ and infinite index set J, such that $|V_k^{n_j}| > \epsilon \cdot n_j$ for $j \in J$. We restrict attention to index set J and consider $\sigma = \{\sigma^n\}_{n=1}^{\infty}$ an equilibrium along which asymptotic learning occurs in the society.

Consider a collection of subsets of the possible realizations of the private signals, $\{Q^n\}_{n=1}^{\infty}$, and a collection of subsets of agents, $\{R^n\}_{n=1}^{\infty}$, such that:

(i) $\lim_{n\to\infty} \mathbb{P}_\sigma(Q^n) > \epsilon$, for some $\epsilon > 0$, i.e., the subsets of realizations considered have positive measure as the society grows, and if a realization is in Q^n its complement is also in Q^n.

(ii) $\lim_{n\to\infty} \frac{1}{n}|R^n| > \epsilon$.

Such collections should exist, since condition (6.2) fails to hold in the society. Consider next equilibrium σ and assume that the realization of the private signals belongs to subset Q^n (for the appropriate n). Since asymptotic learning occurs along equilibrium σ we have:

$$\lim_{n\to\infty} \frac{1}{n}|R_\sigma^n| = 0,$$

where $R_\sigma^n = \{i \in R^n \big| \sigma_{i,d_i^n}^n \in \{0,1\}\}$. However, this implies that there should exist a collection of subsets of agents, $\{S^n\}_{n=1}^{\infty}$, such that:

(i) $R_\sigma^{n,c} \subseteq S_{follow}^n$, where $R_\sigma^{n,c} = \{i \in R^n \big| \sigma_{i,d_i^n}^n = \text{“}wait\text{”}\}$.

(ii) $\sigma_{j,\tau}^n \neq \{\text{“}wait\text{''}\}$, for some $\tau < \tau_i^n - dist(i,j)$ and i such that $j \in B_{i,\tau_i^n}^n$, $i \in R_\sigma^n$.

(iii) $\lim_{n\to\infty} \frac{1}{n}|S^n| = 0$, since otherwise asymptotic learning would not occur along equilibrium σ.

Note that collection $\{S^n\}_{n=1}^{\infty}$ satisfies the definition of a set of leaders [cf. Definition 5],

since

$$\limsup_{n\to\infty} \frac{1}{n}|R^{n,c}| > \epsilon$$

and Proposition 16 (ii) follows. ■

Proof of Theorem 1.

The proof of the theorem relies heavily on the next proposition, which intuitively states that there is no incentive to lie to an agent with a large number of neighbors, assuming that everybody else is truthful.

Proposition 27 (Truthful Communication to a High Degree Agent). *There exists a scalar $k > 0$, such that truth-telling to agent i, with $indeg_i^n \geq k(\bar{\epsilon})$, in the first time period is an equilibrium of $INFO(G^n)$. Formally,*

$$(\sigma^{n,truth}, m^{n,truth}) \in INFO(G^n),$$

where $m_{ji,0}^{n,truth} = s_j$ for $j \in B_{i,1}^n$.

Proof. The proof is based on the following argument. Suppose that all agents in $B_{i,1}^n$ except j report their signals truthfully to i. Moreover, let $|B_{i,1}^n| \geq k$, where k is a large constant. Then, it is an weakly dominant strategy for j to report her signal truthfully to i, since j's message will not be *pivotal* for agent i, i.e., i will take an irreversible action after the first communication step, no matter what j reports. ■

Theorem 1 follows directly from Proposition 27 and Proposition 16. ■

Proof of Theorem 2. The proof uses similar arguments with the ones presented in the proof of Proposition 26. A key observation is that a social planner cannot achieve a higher aggregate welfare than that achieved by the agents under Assumption 5. ■

Proof of Proposition 17. The claim follows by noting that the social planner could choose the following strategy profile: for each $j \in D_{k,\ell}^{n,\sigma}$ delay i's irreversible action by at least one time period, where i is an agent such that if i delays then j gains access to a least

ℓ additional signals. Moreover, it is straightforward to see that there exist ϵ, δ for which ϵ, δ-learning fails. ■

Appendix E

Omitted Proofs from Chapter 7

Proof of Proposition 18

(i) Let the discount rate be $r > 0$. Then, the perfect observation radius of an agent is bounded by a constant Z and, thus, $|V_{k^Z}^n| = n$ for all $n > k^Z$. Thus, perfect asymptotic learning fails from Proposition 16.

(ii) First, we state a lemma from [16] for bounding the probability that there exists a path between two agents in a preferential attachment network. For the remainder, of the proof, agent i is simply the agent that was born at time i.

Lemma in [16] Let $\{G^n\}_{n=1}^{\infty}$ be a preferential attachment sequence of communication networks and consider network G^n, for n large enough. Let P be a path from agent i to agent j in G^n, where $\mathcal{E}^n(P)$ denotes the set of edges of path P. Then,

$$\mathbb{P}(P \subseteq G^n) \leq C^{|\mathcal{E}^n(P)|} \prod_{(i_1,i_2)\in\mathcal{E}^n(P)} \frac{1}{\sqrt{i_1 i_2}},$$

where C is an absolute constant.

This lemma allows us to consider a random communication network, where an edge between agents i and j is present independently with probability $1/\sqrt{ij}$, although in the original network G^n edges are dependent. The intuition why such a result is true comes from

observing that the dependencies between edges due to the preferential attachment rule are not too strong as the communication network grows.

We are now ready to proceed with the proof. Let w^n be the average degree for G^n, which by the preferential attachment rule remains finite with high probability as n grows, i.e., $\lim_{n\to\infty} w^n < W$. In particular, for sufficiently large n,

$$\mathbb{P}(\text{at most } n/c \text{ nodes have degree more than } cW) > 1 - \epsilon/2, \text{ for all } c > 0. \tag{E.1}$$

Next, we use the above lemma to show that most low degree agents agents are not connected with a short path to high degree agents. In particular, let S_1 (S_2) denote the set of agents that have degree greater (smaller) than cW. We want to bound the probability that there exists a short path (at most L links long) from an agent in S_2 to an agent in S_1. As mentioned above, with high probability $|S_1| \leq n/c$. Then, if P_k^ℓ denotes a path of length ℓ from agent $k \in S_2$ to any agent in S_1, we have

$$\begin{aligned}
\sum_{\ell=1}^{L} \mathbb{P}(P_k^\ell \subseteq G^n) &\leq \sum_{\ell=1}^{L} \sum_{j \in S_1} \frac{C^\ell}{\sqrt{k \cdot j}} \sum_{u_1,\cdots,u_{\ell-1} \in S_2} \prod_{t=1}^{\ell-1} \frac{1}{u_t} && \text{(E.2)}\\
&\leq \sum_{\ell=1}^{L} \sum_{j \in S_1} \frac{C^\ell}{\sqrt{k \cdot j}} \sum_{i \in S_2} \frac{1}{i} \\
&\leq \sum_{\ell=1}^{L} \sum_{j \in \{1,\cdots,n/c\}} \frac{C^\ell}{\sqrt{k \cdot j}} \sum_{i \in \{n/c,\cdots,n\}} \frac{1}{i} && \text{(E.3)}\\
&\leq \zeta < 1, \text{ for agents with } k = O(n), && \text{(E.4)}
\end{aligned}$$

where Eq. (E.2) follows from Lemma in [16] and Eq. (E.3) follows by setting set S_1 to be the set of the first n/c nodes (lowest indices) and set S_2 the rest. Finally, Eq. (E.4) follows for an appropriate choice of constant c (and consequently the size of set S_1). In particular, as n grows, we obtain $\sum_{i \in \{n/c,\cdots,n\}} \frac{1}{i} \approx \log c$ and $\sum_{j \in \{1,\cdots,n/c\}} \frac{C^\ell}{\sqrt{k \cdot j}} \approx \frac{constant}{c^{1/2}}$ for $k = O(n)$. Finally, note that with high probability most of the agents with index $O(n)$ belong to set

S_2, since by Lemma in [16] above their expected degree is bounded by a small constant. Combining Eqs. (E.1), (E.4) and Corollary 2 we conclude that perfect asymptotic learning fails in a preferential attachment communication network sequence with probability at least $1-\epsilon$ for discount rate $r > 0$.

(iii) Follows directly from (ii), since Erdős-Renyi communication networks are even sparser than preferential attachment communication networks.

Proof of Proposition 19

(i) - (ii) The claims follow directly from Proposition 16. In particular, for $r < \bar{r}_1$ (for complete) and $r < \bar{r}_2$ (for star), where $\bar{r}_1, \bar{r}_2$ such that:

$$\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2} < \pi \frac{\lambda}{\lambda + \bar{r}_1} \text{ and } \pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2} < \pi \left(\frac{\lambda}{\lambda + \bar{r}_2} \right)^2$$

we obtain that

$$\lim_{k \to \infty} \lim_{n \to \infty} \frac{1}{n} \cdot \left| V_k^n \right| = 0,$$

therefore perfect asymptotic learning occurs in the respective societies.

(iii) Follows directly from [70].

(iv) Consider the following two events A and B.

Event A: Layer 1 (the top layer) has more than k agents, where $k > 0$ is a scalar.

Event B: The total number of layers is more than k.

From the definition of a hierarchical sequence of communication networks, we have

$$\mathbb{P}(A) = \prod_{i=2}^{k} \left(1 - \frac{1}{i^{1+\zeta}} \right) < \exp\left(- \sum_{i=2}^{k} \frac{1}{i^{1+\zeta}} \right). \tag{E.5}$$

Also,

$$\mathbb{P}(B) \leq \frac{\mathbb{E}(\mathcal{L})}{k} = \frac{1}{k} \sum_{i=2}^{\infty} \frac{1}{i^{1+\zeta}}, \tag{E.6}$$

from Markov's inequality, where $\mathcal{L}$ is a random variable, that denotes the number of layers in the hierarchical society. Let $\zeta(\epsilon)$ be small enough and k (and consequently n) large enough such that $\sum_{i=2}^{k} \frac{1}{i^{1+\zeta}} > \log \frac{4}{\epsilon}$ and $\sum_{i=2}^{\infty} \frac{1}{i^{1+\zeta}} < \frac{k \cdot \epsilon}{4}$. For those values of ζ and k we obtain $\mathbb{P}(A) < \epsilon/4$ and $\mathbb{P}(B) < \epsilon/4$. Next, consider the event $C = A^c \cap B^c$, which from Eqs. (E.5) and (E.6) has probability $\mathbb{P}(C) > 1 - \epsilon/2$ for the values of ζ and k chosen above. Moreover, we consider

Event D: The agents on the top layer are information hubs, i.e.,

$$\lim_{n \to \infty} |B_{i,1}^n| = \infty, \text{ for all } i \in \mathcal{N}_1^n.$$

We claim that event D occurs with high probability if event C occurs, i.e., $\mathbb{P}(D| \, C) > 1 - \epsilon/2$, which implies

$$\mathbb{P}(C \cap D) = \mathbb{P}(D| \, C)\mathbb{P}(C) > (1 - \epsilon/2)^2 > 1 - \epsilon. \tag{E.7}$$

In particular, note that conditional on event C occurring, the total number of layers and the total number of agents in the top layer is at most k. From the definition of a hierarchical society, agents in layers with index $\ell > 1$ have an edge to a uniform agent that belongs to a layer with lower index, with probability one. Therefore, if we denote the degree of an agent in a top layer (say agent 1) by $\mathcal{D}_1^n$ we have

$$\mathcal{D}_1^n = \sum_{i=1}^{\mathcal{T}_2^n} \mathcal{I}_{i,1}^{level2} + \cdots + \sum_{i=1}^{\mathcal{T}_{\mathcal{L}}^n} \mathcal{I}_{i,1}^{level\mathcal{L}}, \tag{E.8}$$

where $\mathcal{T}_i^n$ denotes the random number of agents in layer i and $\mathcal{I}_{i,1}^{levelj}$ is an indicator variable that takes value one if there is an edge from agent i to agent 1 (here $levelj$ simply denotes that agent i belongs to level j). Again from the definition, we have $\mathbb{P}(I_{i1}^{levelj} = 1) = \frac{1}{\sum_{\ell=1}^{j-1} \mathcal{T}_\ell^n}$, where the sum in the denominator is simply the total number of agents that lie in layers with lower index, and finally, $\mathcal{T}_1^n + \cdots \mathcal{T}_{\mathcal{L}}^n = n$.

We can obtain a lower bound on the expected degree of an agent in the top layer conditional on event C by viewing (E.8) as the following optimization problem:

$$\begin{aligned} \min \quad & \frac{x_2}{x_1} + \cdots + \frac{x_k}{x_1 + \cdots + x_{k-1}} \\ s.t. \quad & \textstyle\sum_{j=1}^{k} x_j = n, \\ & 0 \leq x_1 \leq k, \\ & 0 \leq x_2, \cdots, x_{k-1}, \end{aligned}$$

where we make use of the fact that the total number of layers is bounded by k, since we condition on event C. By solving the problem we obtain that the objective function is lower bounded by $\phi(n)$, where $\phi(n) = O(n^{1/k})$ for every n. Then,

$$\begin{aligned} &\mathbb{E}[\mathcal{D}_1^n | C] = \\ &= \sum_{\ell=2}^{k} \sum_{\substack{k_1 \leq k, \cdots, k_\ell \\ k_1 + \cdots + k_\ell = n}} \mathbb{P}(\mathcal{L} = \ell, \mathcal{T}_1^n = k_1, \cdots, \mathcal{T}_\ell^n = k_\ell | C) \cdot \mathbb{E}[\mathcal{D}_1^n | C, \mathcal{L} = \ell, \mathcal{T}_1^n = k_1, \cdots, \mathcal{T}_\ell^n = k_\ell] \\ &\geq \sum_{\ell=2}^{k} \sum_{\substack{k_1 \leq k, \cdots, k_\ell \\ k_1 + \cdots + k_\ell = n}} \mathbb{P}(\mathcal{L} = \ell, \mathcal{T}_1^n = k_1, \cdots, \mathcal{T}_\ell^n = k_\ell | C) \cdot \phi(n) = \phi(n), \end{aligned} \tag{E.9}$$

where Eq. (E.9) follows since $\mathbb{E}[\mathcal{D}_1^n | C, \mathcal{L} = \ell, \mathcal{T}_1^n = k_1, \cdots, \mathcal{T}_\ell^n = k_\ell] \geq \phi(n)$ for all values of ℓ $(2 \leq \ell \leq k)$ and $k_1, \cdots, k_\ell$ $(k_1 \leq k, k_1 + \cdots + k_\ell = n)$ from the optimal solution of the optimization problem. The same lower bound applies for all agents in the top layer. Similarly we have for the variance of the degree of an agent in the top layer (we use

$\ell, k_1, \cdots, k_\ell$ as a shorthand for $\mathcal{L} = \ell, \mathcal{T}_1^n = k_1, \cdots, \mathcal{T}_\ell^n = k_\ell$)

$$
\begin{aligned}
Var[\mathcal{D}_1^n|C] &= \sum_{\ell=2}^{k} \sum_{\substack{k_1 \leq k, \cdots, k_\ell \\ k_1+\cdots+k_\ell = n}} \mathbb{P}(\ell, k_1, \cdots, k_\ell|C) \cdot Var[\mathcal{D}_1^n|C, \ell, k_1, \cdots, k_\ell] \\
&= \sum_{\ell=1}^{k} \sum_{\substack{k_1 \leq k, \cdots, k_\ell \\ k_1+\cdots+k_\ell = n}} \mathbb{P}(\ell, k_1, \cdots, k_\ell|C) \cdot \left(k_2 Var(I_{i,1}^{level2}) + \cdots + k_\ell Var(I_{i,1}^{level\ell})\right) \qquad \text{(E.10)} \\
&\leq \sum_{\ell=1}^{k} \sum_{\substack{k_1 \leq k, \cdots, k_\ell \\ k_1+\cdots+k_\ell = n}} \mathbb{P}(\ell, k_1, \cdots, k_\ell|C) \cdot \left(k_2 \mathbb{E}(I_{i,1}^{level2}) + \cdots + k_\ell \mathbb{E}(I_{i,1}^{level\ell})\right) = \mathbb{E}[\mathcal{D}_1^n|C], \qquad \text{(E.11)}
\end{aligned}
$$

where Eq. (E.10) follows by noting that conditional on event C and the number of layers and the agents in each layer being fixed, the indicator variables (defined above) are independent and Eq. (E.11) follows since the variance of an indicator variable is smaller that its expectation. We conclude that the variance of the degree is smaller than the expected value and from Chebyschev's inequality we conclude that

$$\mathbb{P}(D) \geq \mathbb{P}\Big(\bigcap_{i \in \mathcal{N}_1^n} \frac{\mathcal{D}_i^n}{\phi(n)} > \zeta\Big) > 1 - \epsilon/2,$$

where $\zeta > 0$, i.e., with high probability all agents in the top layer are information hubs (recall that $\lim_{n\to\infty} \phi(n) = \infty$).

We have shown that when event $C \cap D$ occurs, there is a path of length at most k (the total number of layers) from each agent to an agent at the top layer, i.e., an information hub with high probability. Therefore, if the discount rate r is greater than some lower bound ($r > \bar{r}_4$), then perfect asymptotic learning occurs. Finally, we complete the proof by noting that $\mathbb{P}(C \cap D) > (1 - \epsilon/2)^2 > 1 - \epsilon$.

Appendix F

Omitted Proofs from Chapter 8

Proof of Theorem 3

First we make an observation which will be used frequently in the subsequent analysis. Consider an agent i such that $H^n_{sc(i)} \in \mathcal{H}^n_{\bar{k}}$, where $\bar{k}$ is an integer appropriately chosen (see below), i.e., the size of the social clique of agent i is greater than or equal to $\bar{k}$, $|H^n_{sc(i)}| \geq \bar{k}$. Suppose agent i does not form a link with cost c with any agents outside her social clique. If she makes a decision at time $t = 0$ based on her signal only, her expected payoff will be $\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}$. If she waits for one period, she has access to signals of all the agents in her social clique (i.e., she has access to at least $\bar{k}$ signals), implying that her expected payoff would be bounded from below by $\frac{\lambda}{r+\lambda}\left(\pi - \frac{1}{1/\rho^2+1/\bar{\rho}^2\bar{k}}\right)$. Hence, her expected payoff $\mathbb{E}[\Pi_i(g^n)]$ satisfies

$$\mathbb{E}[\Pi_i(g^n)] \geq \max\left\{\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}, \frac{\lambda}{r+\lambda}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2\bar{k}}\right)\right\},$$

for any link formation strategy g^n and along any $\sigma \in INFO(G^n)$ (where G^n is the communication network induced by g^n).

Suppose now that agent i forms a link with cost c with an agent outside her social

clique. Then, her expected payoff will be bounded from above by

$$\mathbb{E}[\Pi_i(g^n)] < \max\left\{\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}, \left(\frac{\lambda}{\lambda + r}\right)^2 \pi - c\right\},$$

where the second term in the maximum is an upper bound on the payoff she could get by having access to signals of all agents she is connected to in two time steps (i.e., signals of the agents in her social clique and in the social clique that she is connected to). Combining the preceding two relations, we see that an agent i with $H^n_{sc(i)} \in \mathcal{H}^n_{\bar{k}}$ *will not* form any costly links in any network equilibrium, i.e.,

$$g^n_{ij} = 1 \text{ if and only if } sc(j) = sc(i) \text{ for all } i \text{ such that } |H^n_{sc(i)}| \geq \bar{k}. \tag{F.1}$$

for $\bar{k}$ such that

$$\frac{\lambda}{r + \lambda}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2\bar{k}}\right) \geq \left(\frac{\lambda}{\lambda + r}\right)^2 \pi - c.$$

(a) Condition (8.2) implies that for all sufficiently large n, we have

$$\left|\mathcal{H}^n_{\bar{k}}\right| \geq \xi n, \tag{F.2}$$

where $\xi > 0$ is a constant. For any ϵ with $0 < \epsilon < \xi$, we have

$$\begin{aligned}
\mathbb{P}\left(\sum_{i=1}^{n} \frac{1 - M_i^{n,\epsilon}}{n} > \epsilon\right) &= \mathbb{P}\left(\left[\sum_{i \mid |H^n_{sc(i)}| < \bar{k}} \frac{1 - M_i^{n,\epsilon}}{n} + \sum_{i \mid |H^n_{sc(i)}| \geq \bar{k}} \frac{1 - M_i^{n,\epsilon}}{n}\right] > \epsilon\right) \\
&\geq \mathbb{P}\left(\sum_{i \mid |H^n_{sc(i)}| \geq \bar{k}} \frac{1 - M_i^{n,\epsilon}}{n} > \epsilon\right).
\end{aligned} \tag{F.3}$$

The right-hand side of the preceding inequality can be re-written as

$$\mathbb{P}\left(\sum_{i|\ |H^n_{sc(i)}|\geq \bar{k}} \frac{1-M_i^{n,\epsilon}}{n} > \epsilon\right) = 1- \ \mathbb{P}\left(\sum_{i|\ |H^n_{sc(i)}|\geq \bar{k}} \frac{1-M_i^{n,\epsilon}}{n} \leq \epsilon\right)$$

$$= 1- \ \mathbb{P}\left(\sum_{i|\ |H^n_{sc(i)}|\geq \bar{k}} \frac{M_i^{n,\epsilon}}{n} \geq r-\epsilon\right),$$

where $r = \sum_{i|\ |H^n_{sc(i)}|\geq \bar{k}} \frac{1}{n}$. By Eq. (F.2), it follows that for n sufficiently large, we have $r \geq \xi$. Using Markov's inequality, the preceding relation implies

$$\mathbb{P}\left(\sum_{i|\ |H^n_{sc(i)}|\geq \bar{k}} \frac{1-M_i^{n,\epsilon}}{n} > \epsilon\right) \geq 1 - \frac{\sum_{i|\ |H^n_{sc(i)}|\geq \bar{k}} \mathbb{E}[M_i^{n,\epsilon}]}{n} \cdot \frac{1}{r-\epsilon}. \tag{F.4}$$

By Lemma 8 and observation (F.1), $\mathbb{E}[M_i^{n,\epsilon}]$ for an agent i with $|H^n_{sc(i)}| \geq \bar{k}$ is upper bounded by

$$\mathbb{P}(M_i^{n,\epsilon}=0) \geq erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2|H^n_{sc(i)}|}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2|H^n_{sc(i)}|}}\right),$$

and therefore

$$\mathbb{E}[M_i^{n,\epsilon}] \leq 1 - \left(erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2|H^n_{sc(i)}|}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2|H^n_{sc(i)}|}}\right)\right).$$

Using this bound and assuming without loss of generality, that social cliques are ordered

by size (H_1^n is the biggest), we can re-write Eq. (F.4) as

$$\mathbb{P}\left(\sum_{i|\ |H^n_{sc(i)}|\geq \bar{k}} \frac{1-M_i^{n,\epsilon}}{n} > \epsilon\right) \geq$$
$$\geq 1 - \frac{\sum_{j=1}^{|\mathcal{H}^n_{\bar{k}}|} |H_j^n| \left(1 - \left(erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2 |H_j^n|}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2 |H_j^n|}}\right)\right)\right)}{(r-\epsilon)\cdot n}$$
$$\geq 1 - \frac{r\cdot(1-\zeta)}{r-\epsilon} \geq 1 - \frac{\xi\cdot(1-\zeta)}{\xi-\epsilon} > \delta \tag{F.5}$$

Here, the second inequality is obtained since the largest value for the sum is achieved when all summands are equal and

$$\zeta = erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2\bar{k}}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2\bar{k}}}\right).$$

The third inequality holds using the relation $r \geq \xi$ and choosing appropriate values for ϵ, δ.

This establishes that for all sufficiently large n, we have

$$\mathbb{P}\left(\sum_{i=1}^{n} \frac{1-M_i^n}{n} > \epsilon\right) > \delta > 0,$$

which implies

$$\limsup_{n\to\infty} \mathbb{P}\left(\sum_{i=1}^{n} \frac{1-M_i^n}{n} > \epsilon\right) > \delta,$$

and shows that perfect asymptotic learning does not occur in any network equilibrium.

(b) We show that if the communication cost structure satisfies condition (8.3), then asymptotic learning occurs in all network equilibria $(g,\sigma) = (\{g^n,\sigma^n\})_{n=1}^{\infty}$. For an illustration of the resulting communication networks, when condition (8.4) holds, refer to Figure 8-2(a). Let $B_i^n(G^n)$ be the *neighborhood* of agent i in communication network G^n

(induced by the link formation strategy g^n),

$$B_i^n(G^n) = \{j \mid \text{ there exists a path } \mathcal{P} \text{ in } G^n \text{ from } j \text{ to } i\},$$

i.e., $B_i^n(G^n)$ is the set of agents in G^n whose information agent i can acquire over a sufficiently large (but finite) period of time.

We first show that for any agent i such that $\limsup_{n\to\infty} |H_{sc(i)}^n| < \bar{k}$, her neighborhood in any network equilibrium satisfies $\lim_{n\to\infty} |B_i^n| = \infty$. We use the notion of an isolated social clique to show this. For a given n, we say that a social clique H_ℓ^n is *isolated* (at a network equilibrium (g, σ)) if no agent in H_ℓ^n forms a costly link with an agent outside H_ℓ^n in (g, σ). Equivalently, a social clique H_ℓ^n is not isolated if there exists at least one agent $j \in H_\ell^n$, such that j incurs cost c and forms a link with an agent outside H_ℓ^n.

We show that for an agent i with $\limsup_{n\to\infty} |H_{sc(i)}^n| < \bar{k}$, the social clique $H_{sc(i)}^n$ is not isolated in any network equilibrium for all sufficiently large n. Using condition (8.3), we can assume without loss of generality that social cliques are ordered by size from largest to smallest and that $\lim_{n\to\infty} |H_1^n| = \infty$. Suppose that $H_{sc(i)}^n$ is isolated in a network equilibrium (g, σ). Then the expected payoff of agent i is upper bounded (similarly with above)

$$\mathbb{E}[\Pi_i(g^n)] \leq \max\left\{\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}, \frac{\lambda}{r+\lambda}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2(\bar{k}-1)}\right)\right\}$$

Using the definition of $\bar{k}$, it follows that for some $\epsilon > 0$,

$$\mathbb{E}[\Pi_i(g^n)] \leq \max\left\{\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}, \left(\frac{\lambda}{r+\lambda}\right)^2 \pi - c - \epsilon\right\} \tag{F.6}$$

Suppose next that agent i forms a link with an agent $j \in H_1^n$. Her expected payoff

$\mathbb{E}[\Pi_i(g^n)]$ satisfies

$$\mathbb{E}[\Pi_i(g^n)] \geq \left(\frac{\lambda}{r+\lambda}\right)^2 \cdot \left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 \left|H_1^n\right|}\right) - c,$$

since in two time steps, she has access to the signals of all agents in the social clique H_1^n. Since $\lim_{n\to\infty} |H_1^n| = \infty$, there exists some N_1 such that

$$\mathbb{E}[\Pi_i(g^n)] > \left(\frac{\lambda}{\lambda+r}\right)^2 \pi - c - \epsilon \qquad \text{for all } n > N_1.$$

Comparing this relation with Eq. (F.6), we conclude that under the assumption that $r < \bar{r}$ (for appropriate $\bar{r}$), the social clique $H^n_{sc(i)}$ is not isolated in any network equilibrium for all $n > N_1$.

Next, we show that $\lim_{n\to\infty} |B_i^n| = \infty$ in any network equilibrium. Assume to arrive at a contradiction that $\limsup_{n\to\infty} |B_i^n| < \infty$ in some network equilibrium. This implies that $\limsup_{n\to\infty} |B_i^n| < |H_1^n|$ for all $n > N_2 > N_1$. Consider some $n > N_2$. Since $H^n_{sc(i)}$ is not isolated, there exists some $j \in H^n_{sc(i)}$ such that j forms a link with an agent h outside $H^n_{sc(i)}$. Since $\limsup_{n\to\infty} |B_i^n| < |H_1^n|$, agent j can improve her payoff by changing her strategy to $g^n_{jh} = 0$ and $g^n_{jh'} = 1$ for $h' \in H_1^n$, i.e., j is better off deleting her existing costly link and forming one with an agent in social clique H_1^n. Hence, for any network equilibrium, we have

$$\lim_{n\to\infty} |B_i^n| = \infty \quad \text{for all } i \text{ with } \limsup_{n\to\infty} |H^n_{sc(i)}| < \bar{k} \tag{F.7}$$

We next consider the probability that a non-negligible fraction (ϵ-fraction) of agents takes an action that is at least ϵ-away from optimal with probability at least δ along a network equilibrium (g, σ). For any n, we have from Markov's inequality

$$\mathbb{P}\left(\sum_{i=1}^{n} \frac{1 - M_i^{n,\epsilon}}{n} > \epsilon\right) \leq \frac{1}{\epsilon} \cdot \sum_{i=1}^{n} \frac{\mathbb{E}[1 - M_i^{n,\epsilon}]}{n} \tag{F.8}$$

We next provide upper bounds on the individual terms in the sum on the right-hand side. For any agent i, we have

$$\mathbb{E}[1 - M_i^{n,\epsilon}] \leq erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2|B_i^n|}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2|B_i^n|}}\right). \tag{F.9}$$

Consider an agent i with $\limsup_{n\to\infty} |H_{sc(i)}^n| < \bar{k}$ (i.e., $|H_{sc(i)}^n| < \bar{k}$ for all n large). By Eq. (F.7), we have $\lim_{n\to\infty} |B_i^n| = \infty$. Together with Eq. (F.9), this implies that for some $\zeta > 0$, there exists some N such that for all $n > N$, we have

$$\mathbb{E}[1 - M_i^{n,\epsilon}] < \frac{\epsilon\,\zeta}{2} \qquad \text{for all } i \text{ with } \limsup_{n\to\infty} |H_{sc(i)}^n| < \bar{k}. \tag{F.10}$$

Consider next an agent i with $\limsup_{n\to\infty} |H_{sc(i)}^n| \geq \bar{k}$, and for simplicity, let us assume that the limit exists, i.e., $\lim_{n\to\infty} |H_{sc(i)}^n| \geq \bar{k}$.[1] This implies that $|H_{sc(i)}^n| \geq \bar{k}$ for all large n, and therefore,

$$\begin{aligned}
\sum_{i|\ \limsup_{n\to\infty} |H_{sc(i)}^n| \geq \bar{k}} \frac{\mathbb{E}[1 - M_i^{n,\epsilon}]}{n} &\leq \sum_{j=1}^{|\mathcal{H}_{\bar{k}}^n|} |H_j^n| \cdot \left(erf\left(\frac{\epsilon}{\sqrt{2\bar{\rho}^2|H_j^n|}}\right) - erf\left(-\frac{\epsilon}{\sqrt{2\bar{\rho}^2|H_j^n|}}\right)\right) \\
&\leq \frac{|\mathcal{H}_{\bar{k}}^n|}{n} \cdot \bar{k},
\end{aligned}$$

where the first inequality follows from Eq. (F.9). Using condition (8.3), i.e., $\lim_{n\to\infty} \frac{|\mathcal{H}_{\bar{k}}^n|}{n} = 0$, this relation implies that there exists some $\tilde{N}$ such that for all $n > \tilde{N}$, we have

$$\sum_{i|\ \limsup_{n\to\infty} |H_{sc(i)}^n| \geq \bar{k}} \frac{\mathbb{E}[1 - M_i^{n,\epsilon}]}{n} < \frac{\epsilon\,\zeta}{2}. \tag{F.11}$$

[1]The case when the limit does not exist can be proven by focusing on different subsequences. In particular, along any subsequence $\mathcal{N}_i$ such that $\lim_{n\to\infty, n\in\mathcal{N}_i} |H_{sc(i)}^n| \geq \bar{k}$, the same argument holds. Along any subsequence $\mathcal{N}_i$ with $\lim_{n\to\infty, n\in\mathcal{N}_i} |H_{sc(i)}^n| < \bar{k}$, we can use an argument similar to the previous case to show that $\lim_{n\to\infty, n\in\mathcal{N}_i} |B_i^n| = \infty$, and therefore $\mathbb{E}[1 - M_i^{n,\epsilon}] < \frac{\epsilon\zeta}{2}$ for n large and $n \in \mathcal{N}_i$.

Combining Eqs. (F.10) and (F.11) with Eq. (F.8), we obtain for all $n > \max\{N, \tilde{N}\}$,

$$\mathbb{P}\left(\sum_{i=1}^{n} \frac{1 - M_i^n}{n} > \epsilon\right) < \zeta,$$

where $\zeta > 0$ is an arbitrary scalar. This implies that

$$\lim_{n\to\infty} \mathbb{P}\left(\sum_{i=1}^{n} \frac{1 - M_i^n}{n} > \epsilon\right) = 0,$$

for all ϵ, showing that perfect asymptotic learning occurs along every network equilibrium.

(c) The proof proceeds in two parts. First, we show that if condition (8.4) is satisfied, learning occurs in at least one network equilibrium (g, σ). Then, we show that there exists a $\bar{c} > 0$, such that if $c < \bar{c}$, then learning occurs in all network equilibria. We complete the proof by showing that if $c > \bar{c}$, then there exist network equilibria, in which asymptotic learning fails, even when condition (8.4) holds. We consider the case when agents are patient, i.e., the discount rate $r \to 0$. We consider $\bar{k}$, such that $c > \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 \bar{k}}$ and $c < \frac{1}{1/\rho^2 + 1/\bar{\rho}^2(\bar{k}-1)} - \epsilon'$, for some $\epsilon' > 0$ (such a $\bar{k}$ exists). Finally, we assume that $c < \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}$, since otherwise no agent would have an incentive to form a costly link.

Part 1: We assume, without loss of generality, that social cliques are ordered by size (H_1^n is the smallest). Let $\mathcal{H}_{<\bar{k}}^n$ denote the set of social cliques of size less than $\bar{k}$, i.e., $\mathcal{H}_{<\bar{k}}^n = \{H_i^n, i = 1, \ldots, K^n \mid |H_i^n| < \bar{k}\}$. Finally, let *rec*$(j)$ and *send*(j) denote two special nodes for social clique H_j^n, the *receiver* and the *sender* (they might be the same node). We claim that (g^n, σ^n) described below and depicted in Figure 8-2(b) is an equilibrium of the

network learning game $\Gamma(C^n)$ for n large enough and δ sufficiently close to one.

$$g_{ij}^n = \begin{cases} 1 & \text{if } sc(i) = sc(j), \text{ i.e., } i, j \text{ belong to the same social clique,} \\ 1 & \text{if } i = rec(\ell - 1) \text{ and } j = send(\ell) \text{ for } 1 < \ell \leq |\mathcal{H}^n_{<\bar{k}}|, \\ 1 & \text{if } i = rec(|\mathcal{H}^n_{<\bar{k}}|) \text{ and } j = send(1), \\ 0 & \text{otherwise} \end{cases}$$

and $\sigma^n \in INFO(G^n)$, where G^n is the communication network induced by g^n. In this communication network, social cliques with size less than $\bar{k}$ are organized in a directed *ring*, and all agents i, such that $|H^n_{sc(i)}| < \bar{k}$ have the same neighborhood, i.e., $B_i^n = B^n$ for all such agents.

Next, we show that the strategy profile (g^n, σ^n) described above is indeed an equilibrium of the network learning game $\Gamma(C^n)$. We restrict attention to large enough n's. In particular, let N be such that $\sum_{i=1}^{|\mathcal{H}^N_{<\bar{k}}|} |H_i^N| > \bar{k}$ and consider any $n > N$ (such N exists from condition (8.4)). Moreover, we assume that the discount rate is sufficiently close to zero. We consider the following two cases.

Case 1: Agent i is not a connector. Then, $g_{ij}^n = 1$ if and only if $sc(j) = sc(i)$. Agent i's neighborhood as noted above is set B^n, which is such that $\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 |B^n|} > \pi - c$ from the assumption on n, i.e., $n > N$, where N such that $\sum_{i=1}^{|\mathcal{H}^N_{<\bar{k}}|} |H_i^N| > \bar{k}$. Agent i can communicate with all agents in B^n in at most $|\mathcal{H}_{<\bar{k}}|$ communication steps. Therefore, her expected payoff is lower-bounded by

$$\mathbb{E}[\Pi_i(g^n)] \geq \left(\frac{\lambda}{\lambda + r}\right)^{|\mathcal{H}^n_{<\bar{k}}|} \cdot \left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 \bar{k}}\right) > \pi - c,$$

under any equilibrium σ^n for r sufficiently close to zero. Agent i can deviate by forming a costly link with agent m, such that $sc(m) \neq sc(i)$. However, this is not profitable since from above her expected payoff under (g^n, σ^n) is at least $\pi - c$ (which is the maximum possible payoff if an agent chooses to form a costly link).

Case 2: Agent i is a connector, i.e., there exists exactly one j, such that $sc(j) \neq sc(i)$ and $g_{ij}^n = 1$. Using a similar argument as above we can show that it is not profitable for agent i to form an additional costly link with an agent m, such that $sc(m) \neq sc(i)$. On the other hand, agent i could deviate by setting $g_{ij}^n = 0$. However, then her expected payoff would be

$$\begin{aligned}\mathbb{E}[\Pi_i(g^n)] &= \max\left\{\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}, \frac{\lambda}{r+\lambda}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2|H_i^n|}\right)\right\} \\ &\leq \max\left\{\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}, \frac{\lambda}{r+\lambda}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2(\bar{k}-1)}\right)\right\} < \pi - c - \epsilon' \\ &< \left(\frac{\lambda}{r+\lambda}\right)^{|\mathcal{H}_{<\bar{k}}^n|}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2|B^n|}\right) - c - \epsilon,\end{aligned} \tag{F.12}$$

for discount rate sufficiently close to zero. Therefore deleting the costly link is not a profitable deviation. Similarly we can show that it a (weakly) dominant strategy for the connector not to replace her costly link with another costly link.

We showed that (g^n, σ^n) is an equilibrium of the network learning game. Note that we described a link formation strategy, in which social cliques connect to each other in a specific order (in increasing size). There is nothing special about this ordering and any permutation of the first $|\mathcal{H}_{<\bar{k}}^n|$ cliques is an equilibrium as long as they form a directed ring. Finally, any node in a social clique can be a receiver or a sender.

Next, we argue that asymptotic learning occurs in network equilibria $(g, \sigma) = \{(g^n, \sigma^n)\}_{n=1}^{\infty}$, where for all $n > N$, N is a large constant, g^n has the form described above. As shown above, all agents i for which $H_{sc(i)}^n < \bar{k}$ have the same neighborhood, which we denoted by B^n. Moreover, $\lim_{n\to\infty} |B^n| = \infty$, since all social cliques with size less than $\bar{k}$ are connected to the ring and, by condition (8.4), $\lim_{n\to\infty} \sum_{i|\ |H_i^n|<\bar{k}} |H_i^n| = \infty$. For discount rate r sufficiently close to zero and from arguments similar to those in the proof of part (b), we conclude that asymptotic learning occurs in network equilibria (g, σ).

Part 2: We have shown a particular form of network equilibria, in which asymptotic learn-

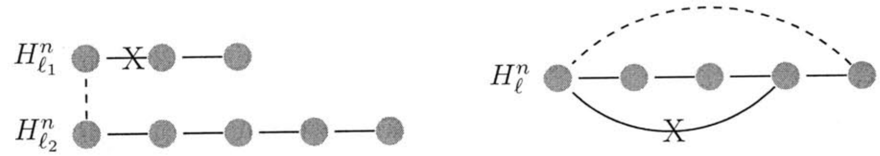

(a) Deviation for $i \in H^n_{\ell_1}$ - property (i). (b) Deviation for $i \in H^n_\ell$ - property (ii).

Figure F-1: Communication networks under condition (8.4).

ing occurs. The following proposition states that for discount rate sufficiently close to zero network equilibria fall in one of two forms.

Proposition 28. *Let Assumptions 6, 8 and condition (8.4) hold. Then, an equilibrium* (g^n, σ^n) *of the network learning game* $\Gamma(C^n)$ *can be in one of the following two forms.*

(i) ***(Incomplete) Ring Equilibrium:*** *Social cliques with indices* $\{1, \cdots, j\}$, *where* $j \leq |\mathcal{H}^n_{<\bar{k}}|$, *form a directed ring as described in Part 1 and the rest of the social cliques are isolated. We call those equilibria ring equilibria and, in particular, a ring equilibrium is called complete if* $j = |\mathcal{H}^n_{<\bar{k}}|$, *i.e., if all social cliques with size less than* $\bar{k}$ *are not isolated.*

(ii) ***Directed Line Equilibrium:*** *Social cliques with indices* $\{1, \cdots, j\}$, *where* $j \leq |\mathcal{H}^n_{<\bar{k}}|$, *and clique with index* $|H^n_{K^n}|$ *(the largest clique) form a directed line with the latter being the endpoint. The rest of the social cliques are isolated.*

Proof. Let (g^n, σ^n) be an equilibrium of the network learning game $\Gamma(C^n)$. Monotonicity of the expected payoff as a function of the number of signals observed implies that if clique H^n_ℓ is not isolated, then no clique with index less than ℓ is isolated in the communication network induced by g^n. In particular, let $conn(\ell)$ be the connector of social clique H^n_ℓ and $\mathbb{E}[\Pi_{conn(\ell)}(g^n)]$ be her expected payoff. Consider an agent i such that $sc(i) = \ell' < \ell$ and, for the sake of contradiction, $H^n_{\ell'}$ is isolated in the communication network induced by g^n. Social cliques are ordered by size, therefore, $|H^n_{\ell'}| \leq |H^n_\ell|$. At this point, we use the

monotonicity mentioned above. Consider the expected payoff of agent i:

$$\begin{aligned}\mathbb{E}[\Pi_i(g^n)] &= \max\left\{\pi - \frac{1}{1/\rho^2+1/\bar{\rho}^2}, \frac{\lambda}{\lambda+r}\left(\pi - \frac{1}{1/\rho^2+1/\bar{\rho}^2|H^n_{\ell'}|}\right)\right\} \\ &\leq \max\left\{\pi - \frac{1}{1/\rho^2+1/\bar{\rho}^2}, \frac{\lambda}{\lambda+r}\left(\pi - \frac{1}{1/\rho^2+1/\bar{\rho}^2|H^n_{\ell}|}\right)\right\} < \mathbb{E}[\Pi_{conn(\ell)}(g^n)],\end{aligned} \tag{F.13}$$

where the last inequality follows from the fact that agent $conn(\ell)$ formed a costly link. Consider a deviation, $g_i^{n,deviation}$ for agent i, in which $g_{i,conn(\ell)}^{n,deviation} = 1$ and $g_{ij}^{n,deviation} = g_{ij}^n$, i.e., agent i forms a costly link with agent $conn(\ell)$. Then,

$$\mathbb{E}[\Pi_i(g^{n,deviation})] \geq \frac{\lambda}{\lambda+r}\,\mathbb{E}[\Pi_{conn(\ell)}(g^n)] > \mathbb{E}[\Pi_i(g^n)],$$

from (F.13) and for discount rate sufficiently close to zero. Therefore, social clique $H^n_{\ell'}$ will not be isolated in any network equilibrium (g^n, σ^n).

Next, we show two structural properties that all network equilibria (g^n, σ^n) should satisfy, when the discount rate r is sufficiently close to one. We say that there exists a path $\mathcal{P}$ between social cliques $H^n_{\ell_1}$ and $H^n_{\ell_2}$, if there exists a path between some $i \in H^n_{\ell_1}$ and $j \in H^n_{\ell_2}$. Also, we say that the in-degree (out-degree) of social clique $H^n_{\ell_1}$ is k, if the sum of in-links (out-links) of the nodes in $H^n_{\ell_1}$ is k, i.e., $H^n_{\ell_1}$ has in-degree k if $\sum_{i\in H^n_{\ell_1}}\sum_{j\notin H^n_{\ell_1}} g^n_{ij} = k$.

(i) Let $H^n_{\ell_1}, H^n_{\ell_2}$ be two social cliques that are not isolated. Then, there should exist a directed path $\mathcal{P}$ in G^n induced by g^n between the two social cliques.

(ii) The in-degree and out-degree of each social clique is at most one.

Figure F-1 provides an illustration of why the properties hold for patient agents. In particular, for property (i), let $i = conn(H^n_{\ell_1})$ and $j = conn(H^n_{\ell_2})$ and assume, without loss of generality, that $|B^n_i| \leq |B^n_j|$. Then, for discount rate sufficiently close to zero and from monotonicity of the expected payoff, we conclude that i has an incentive to deviate, delete her costly and form a costly link with agent j. Property (ii) follows due to similar argu-

ments.

From the above, we conclude that the only two potential equilibrium topologies are the (incomplete) ring and the directed line with the largest clique being the endpoint under the assumptions of the proposition. ■

So far we have shown a particular form of network equilibria, that arise under condition (8.4), in which asymptotic learning occurs. We also argued that under condition (8.4) only (incomplete) ring or directed line equilibria can arise for network learning game $\Gamma(C^n)$. In the remainder we show that there exists a bound $\bar{c} > 0$ on the common cost c for forming a link between two social cliques, such that if $c < \bar{c}$ all network equilibria (g, σ) that arise satisfy that g^n is a complete ring equilibrium for all $n > N$, where N is a constant. In those network equilibria asymptotic learning occurs as argued in Part 1. On the other hand, if $c > \bar{c}$ coordination among the social cliques may fail and additional equilibria arise in which asymptotic learning does not occur. Let

$$\bar{c}^n = \min_k \left\{ -\frac{1}{1/\rho^2 + 1/\bar{\rho}^2(\sum_{j=1}^{k} |H_j^n| + |H_{k+1}^n|)} + \frac{1}{1/\rho^2 + 1/\bar{\rho}^2 |H_{k+1}^n|)} \right\} \tag{F.14}$$

where $k_1 \leq k < |H_{<\bar{k}}^n|$ and $\sum_{j=1}^{k_1} |H_j^n| \geq |H_{K^n}^n|$ (size of the largest social clique). Moreover, let

$$\bar{c} = \lim_{n \to \infty} \inf \bar{c}^n.$$

Then,

Proposition 29. *Let Assumptions 6, 8 and condition (8.4) hold. If $c < \bar{c}$ asymptotic learning occurs in all network equilibria (g, σ). Otherwise, there exist equilibria in which asymptotic learning does not occur.*

Proof. Let the common cost c be such that $c < \bar{c}$, where $\bar{c}$ is defined as above, and consider a network equilibrium (g, σ). Let N be a large enough constant and consider the corresponding g^n for $n > N$. We claim that g^n is a complete ring equilibrium for all such n.

Assume for the sake of contradiction that the claim is not true. Then, from Proposition 28, g^n is either an incomplete ring equilibrium or a directed line equilibrium. We consider the former case (the latter case can be shown with similar arguments). There exists an isolated social clique H_ℓ^n, such that $|H_\ell^n| < \bar{k}$ and all cliques with index less than ℓ are not isolated and belong to the incomplete ring. However, from the definition of $\bar{c}$ we obtain that an agent $i \in H_\ell^n$ would have an incentive to connect to the incomplete ring, thus we reach a contradiction. In particular, consider the following link formation strategy for agent i: $g_{im}^{n,deviation} = 1$ for agent $m \in H_{\ell-1}^n$ and $g_{ij}^{n,deviation} = g_{ij}^n$ for $j \neq m$. Then,

$$\begin{aligned}\mathbb{E}[\Pi_i^n(g^{n,deviation})] &\geq \left(\frac{\lambda}{\lambda+r}\right)^{|H^n_{<\bar{k}}|}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2(\sum_{j=1}^{\ell-1}|H_j^n| + |H_\ell^n|)}\right) - c \\ &> \max\left\{\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2}, \frac{\lambda}{\lambda+r}\left(\pi - \frac{1}{1/\rho^2 + 1/\bar{\rho}^2|H_\ell^n|}\right)\right\} = \mathbb{E}[\Pi_i^n(g^n)],\end{aligned}$$

where the strict inequality follows from the definition of $\bar{c}$ for r sufficiently close to zero. Thus, we conclude that if $c < \bar{c}$, g^n is a complete ring for all $n > N$, where N is a large constant, and from Part 1 asymptotic learning occurs in all network equilibria (g, σ).

On the contrary, if $c > \bar{c}$, then there exists an infinite index set W, such that for all n in the (infinite) subsequence, $\{n_w\}_{w \in W}$, there exists a k, such that

$$\frac{1}{1/\rho^2 + 1/\bar{\rho}^2(\sum_{j=1}^{k}|H_j^n| + |H_{k+1}^n|)} - c < \frac{1}{1/\rho^2 + 1/\bar{\rho}^2|H_{k+1}^n|}. \tag{F.15}$$

Moreover, $|H_{k+1}^n| < \bar{k}$ and $\sum_{j=1}^k |H_j^n| \geq |H_{K^n}^n|$. We conclude that for (F.15) to hold it has to be that $\sum_{j=1}^k |H_j^n| < R$, where R is a uniform constant for all n in the subsequence. Consider $(g, \sigma)_{n=1}^\infty$, such that for every n in the subsequence, g^n is such that social cliques with index greater than k (as described above) are isolated and the rest form an incomplete ring or a directed line and $\sigma^n = INFO(G^n)$, where G^n is the communication network induced by g^n. From above, we obtain that for $c > \bar{c}$, (g^n, σ^n) is an equilibrium of the network learning game $\Gamma(C^n)$. Perfect assymptotic learning, however, fails in such an equilibrium,

since for every $i \in N^n$, $|B_i^n| \leq R$, where recall that B_i^n denotes the neighborhood of agent i. ■

Proposition 29 concludes the proof of Theorem 3.

References

[1] D Acemoglu, K Bimpikis, and A Ozdaglar. Dynamics of information exchange in endogenous social networks. *Working paper*, 2010.

[2] D Acemoglu, K Bimpikis, and A Ozdaglar. Experimentation, patents and innovation. *Working paper*, 2010.

[3] D Acemoglu, M Dahleh, I Lobel, and A Ozdaglar. Bayesian learning in social networks. *Working paper*, Jan 2010.

[4] P Aghion and P Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, Jan 1992.

[5] R Albert and A Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, Jan 2002.

[6] G Angeletos and A Pavan. Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142, Jan 2007.

[7] K Arrow. Economic welfare and the allocation of resources for invention. *The Rate and Direction of Inventive Activity: Economic and Social Factors*, Jan 1962.

[8] D Austen-Smith. Information transmission in debate. *American Journal of Political Science*, 34(1):124–152, Jan 1990.

[9] D Austen-Smith and J Banks. Information aggregation, rationality, and the condorcet jury theorem. *American Political Science Review*, 90(1):34–45, Jan 1996.

[10] V Bala and S Goyal. Learning from neighbours. *Review of Economic Studies*, 65(3):595–621, Jan 1998.

[11] V Bala and S Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, Jan 2000.

[12] A Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, Jan 1992.

[13] A Banerjee and D Fudenberg. Word-of-mouth learning. *Games and Economic Behav-*

ior, 46:1–22, Jan 2004.

[14] A Barabási and R Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, Jan 1999.

[15] S Bikhchandani, D Hirshleifer, and I Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, Jan 1992.

[16] B Bollobás and O Riordan. The diameter of a scale-free random graph. *Combinatorica*, 24(1):5–34, Jan 2004.

[17] P Bolton and C Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, Jan 1999.

[18] J Bulow and P Klemperer. Rational frenzies and crashes. *Journal of Political Economy*, 102(1):1–23, Jan 1994.

[19] J Bulow and P Klemperer. The generalized war of attrition. *American Economic Review*, 89(1):175–189, Jan 1999.

[20] C Chamley and D Gale. Information revelation and strategic delay in a model of investment. *Econometrica*, 62(5):1065–105, Jan 1994.

[21] V Crawford and J Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, Jan 1982.

[22] P Dasgupta. Patents, priority and imitation or, the economics of races and waiting games. *Economic Journal*, 98(389):66–80, Jan 1988.

[23] P DeMarzo, D Vayanos, and J Zwiebel. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118(3):909–968, Jan 2003.

[24] D Duffie, S Malamud, and G Manso. Information percolation with equilibrium search dynamics. *forthcoming in Econometrica*, 2010.

[25] R. Durrett. *Random graph dynamics.* Cambridge University Press, 2007.

[26] A Earls. Digital equipment corporation. *Arcadia Publishing, San Francisco*, Jan 2004.

[27] J Farrell and R Gibbons. Cheap talk with two audiences. *American Economic Review*,

79(5):1214–1223, Jan 1989.

[28] G Folland. Real analysis: Modern techniques and their applications. *John Wiley & Sons, New York*, Jan 1984.

[29] D Fudenberg and J Tirole. A theory of exit in duopoly. *Econometrica*, 54(4):943–960, Jan 1986.

[30] D Gale and S Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, Jan 2003.

[31] A Galeotti. One-way flow networks: the role of heterogeneity. *Economic Theory*, 29:163–179, Jan 2006.

[32] A Galeotti, C Ghiglino, and F Squintani. Strategic information transmission in networks. *Working paper*, Jan 2010.

[33] A Galeotti and S Goyal. The law of the few. *forthcoming in the American Economic Review*, Jan 2010.

[34] N Gallini and S Scotchmer. Intellectual property: When is it the best incentive system? *Innovation policy and the economy*, 2:51–77, Jan 2002.

[35] R Gilbert and C Shapiro. Optimal patent length and breadth. *RAND Journal of Economics*, 21(1):106–112, Jan 1990.

[36] M. Gladwell. *The Tipping Point: How little things can make a big difference.* Little Brown, 2000.

[37] B Golub and M Jackson. Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, Jan 2010.

[38] S. Goyal. *Connections: an introduction to the economics of networks.* Princeton University Press, Princeton, New Jersey, 2007.

[39] G Grossman and E Helpman. Quality ladders in the theory of growth. *Review of Economic Studies*, 58(1):43–61, Jan 1991.

[40] F Gul and R Lundholm. Endogenous timing and the clustering of agents' decisions. *Journal of Political Economy*, 103(5):1039–1066, Jan 1995.

[41] J Hagenbach and F Koessler. Strategic communication networks. *Working paper*, Jan 2010.

[42] J Haigh and C Cannings. The n-person war of attrition. *Acta Applicandae Mathematicae*, 14(1-2):59–74, Jan 1989.

[43] T Harford. ?why giant technological leaps aren?'t always the answer. *Financial Times*, July 18, 2009, 2009.

[44] K Hendricks and D Kovenock. Asymmetric information, information externalities, and efficiency: the case of oil exploration. *RAND Journal of Economics*, 20(2):164–182, Jan 1989.

[45] K Hendricks and R Porter. The timing and incidence of exploratory drilling on offshore wildcat tracts. *American Economic Review*, 86(3):388–407, Jan 1996.

[46] K Hendricks, A Weiss, and C Wilson. The war of attrition in continuous time with complete information. *International Economic Review*, 29(4):663–680, Jan 1988.

[47] M. Jackson. *A survey of models of network formation: stability and efficiency. In group formation in Economics; Networks, clubs and coalitions, edited by Gabrielle Demange and Myrna Wooders.* Cambridge University Press, 2004.

[48] M. Jackson. *Social and economic networks.* Princeton University Press, Princeton, New Jersey, 2008.

[49] M Jackson and A Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, Jan 1996.

[50] M Jovanović, F Annexstein, and K Berman. Modeling peer-to-peer network topologies through "small-world" models and power laws. *Telfor 2001*, Jan 2001.

[51] S Kapur. Technological diffusion with social learning. *Journal of Industrial Economics*, 43(2):173–195, Jan 1995.

[52] G Keller, S Rady, and M Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, Jan 2005.

[53] E Kitch. The nature and function of the patent system. *Journal of Law and Economics*,

20(2):265–290, Jan 1977.

[54] P Klemperer. How broad should the scope of patent protection be? *RAND Journal of Economics*, 21(1):113–130, Jan 1990.

[55] M Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, 1(2):226–251, Jan 2004.

[56] J Morgan and P Stocken. Information aggregation in polls. *American Economic Review*, 98(3):864–896, Jan 2008.

[57] M Newman. Scientific collaboration networks. i. network construction and fundamental results. *Physical Review E*, 64(1):1–8, Jan 2001.

[58] J Reinganum. Dynamic games of innovation. *Journal of Economic Theory*, 25(1):21–41, Jan 1981.

[59] P Romer. Endogenous technological change. *Journal of Political Economy*, 98(5), Jan 1990.

[60] S Scotchmer. On the optimality of the patent renewal system. *RAND Journal of Economics*, 30(2):181–196, Jan 1999.

[61] S Scotchmer. Innovation and incentives. *MIT Press, Cambridge, Massachusetts*, Jan 2004.

[62] H Seyed-Allaei, G Bianconi, and M Marsili. Scale-free networks with an exponent less than two. *Physical Review E*, 73(4):1–5, Jan 2006.

[63] L Simon and M Stinchcombe. Extensive form games in continuous time: Pure strategies. *Econometrica*, 57(5):1171–1214, Jan 1989.

[64] J M Smith. The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, 47(1):209–221, Jan 1974.

[65] L Smith and P Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, Jan 2000.

[66] L Smith and P Sørensen. Rational social learning with random sampling. *Working paper*, Jan 2008.

[67] C Taylor. Digging for golden carrots: an analysis of research tournaments. *American Economic Review*, 85(4):872–890, Jan 1995.

[68] J Tirole. The theory of industrial organization. *MIT Press, Cambridge, Massachusetts*, Jan 1988.

[69] Z Toroczkai and K Bassler. Network dynamics: Jamming is limited in scale-free systems. *Nature*, 428(716), Jan 2004.

[70] H van den Esker, R van der Hofstad, and G Hooghiemstra. Distances in random graphs with infinite mean degrees. *Extremes*, 8(3):111–141, Jan 2005.